

(12) **United States Patent**  
**Wong**

(10) **Patent No.:** **US 8,811,390 B2**  
(45) **Date of Patent:** **\*Aug. 19, 2014**

(54) **SYSTEM AND METHOD FOR HIGH SPEED PACKET TRANSMISSION**

(75) Inventor: **Yuen Fai Wong**, San Jose, CA (US)  
(73) Assignee: **Foundry Networks, LLC**, Santa Clara, CA (US)  
(\* ) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 820 days.  
This patent is subject to a terminal disclaimer.

FOREIGN PATENT DOCUMENTS

EP 1380127 A2 1/2004  
JP 2003/289359 A 10/2003

(Continued)

OTHER PUBLICATIONS

Non-Final Office Action for U.S. Appl. No. 11/646,845, mailed on Oct. 14, 2011, 19 pages.

(Continued)

(21) Appl. No.: **12/608,972**  
(22) Filed: **Oct. 29, 2009**

(65) **Prior Publication Data**

US 2010/0061393 A1 Mar. 11, 2010

**Related U.S. Application Data**

(62) Division of application No. 11/000,359, filed on Nov. 29, 2004, now Pat. No. 7,636,369, which is a division of application No. 10/438,545, filed on May 15, 2003, now Pat. No. 6,901,072.

(51) **Int. Cl.**  
**H04L 12/50** (2006.01)  
**H04Q 11/04** (2006.01)  
**H04L 12/56** (2006.01)

(52) **U.S. Cl.**  
CPC ..... **H04Q 11/0428** (2013.01); **H04L 49/30** (2013.01)  
USPC ..... **370/359**; 370/419

(58) **Field of Classification Search**  
CPC ..... H04Q 11/0428; H04L 49/30  
USPC ..... 370/360, 363, 365, 389, 368, 371, 374, 370/381, 395.7, 395.71, 419, 428, 423, 463, 370/474, 359, 396, 395.6  
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

3,866,175 A 2/1975 Seifert, Jr. et al.  
4,325,119 A 4/1982 Grandmaison et al.

(Continued)


Primary Examiner — Brenda H Pham

(74) *Attorney, Agent, or Firm* — Kilpatrick Townsend & Stockton LLP

(57) **ABSTRACT**

The present invention provides systems and methods for providing data transmission speeds at or in excess of 10 gigabits per second between one or more source devices and one or more destination devices. According to one embodiment, the system of the present invention comprises a first and second media access control (MAC) interfaces to facilitate receipt and transmission of packets over an associated set of physical interfaces. The system also contemplates a first and second field programmable gate arrays (FPGA) coupled to the MAC interfaces and an associated first and second memory structures, the first and second FPGAs are configured to perform initial processing of packets received from the first and second MAC interfaces and to schedule the transmission of packets to the first and second MAC interface for transmission to one or more destination devices. The first and second FPGAs are further operative to dispatch and retrieve packets to and from the first and second memory structures. A third FPGA, coupled to the first and second memory structures and a backplane, is operative to retrieve and dispatch packets to and from the first and second memory structures, compute appropriate destinations for packets and organize packets for transmission. The third FPGA is further operative to receive and dispatch packets to and from the backplane.

**10 Claims, 13 Drawing Sheets**





(56)

## References Cited

## U.S. PATENT DOCUMENTS

|           |   |         |                    |           |    |         |                       |
|-----------|---|---------|--------------------|-----------|----|---------|-----------------------|
| 4,348,725 | A | 9/1982  | Farrell et al.     | 5,907,660 | A  | 5/1999  | Inoue et al.          |
| 4,628,480 | A | 12/1986 | Floyd              | 5,909,686 | A  | 6/1999  | Muller et al.         |
| 4,667,323 | A | 5/1987  | Engdahl et al.     | 5,915,094 | A  | 6/1999  | Kouloheris et al.     |
| 4,679,190 | A | 7/1987  | Dias et al.        | 5,920,566 | A  | 7/1999  | Hendel et al.         |
| 4,683,564 | A | 7/1987  | Young et al.       | 5,920,886 | A  | 7/1999  | Feldmeier             |
| 4,698,748 | A | 10/1987 | Juzswik et al.     | 5,936,939 | A  | 8/1999  | Des Jardins et al.    |
| 4,723,243 | A | 2/1988  | Joshi et al.       | 5,936,966 | A  | 8/1999  | Ogawa et al.          |
| 4,754,482 | A | 6/1988  | Weiss              | 5,956,347 | A  | 9/1999  | Slater                |
| 4,791,629 | A | 12/1988 | Burns et al.       | 5,999,528 | A  | 12/1999 | Chow et al.           |
| 4,794,629 | A | 12/1988 | Pastyr et al.      | 6,000,016 | A  | 12/1999 | Curtis et al.         |
| 4,807,280 | A | 2/1989  | Posner et al.      | 6,011,910 | A  | 1/2000  | Chau et al.           |
| 4,876,681 | A | 10/1989 | Hagiwara et al.    | 6,016,310 | A  | 1/2000  | Muller et al.         |
| 4,896,277 | A | 1/1990  | Vercellotti et al. | 6,023,471 | A  | 2/2000  | Haddock et al.        |
| 4,985,889 | A | 1/1991  | Frankish et al.    | 6,031,843 | A  | 2/2000  | Swanbery et al.       |
| 5,101,404 | A | 3/1992  | Kunimoto et al.    | 6,035,414 | A  | 3/2000  | Okazawa et al.        |
| 5,136,584 | A | 8/1992  | Hedlund            | 6,038,288 | A  | 3/2000  | Thomas et al.         |
| 5,195,181 | A | 3/1993  | Bryant et al.      | 6,067,298 | A  | 5/2000  | Shinohara             |
| 5,208,856 | A | 5/1993  | Leduc et al.       | 6,067,606 | A  | 5/2000  | Holscher et al.       |
| 5,224,108 | A | 6/1993  | McDysan et al.     | 6,076,115 | A  | 6/2000  | Sambamurthy et al.    |
| 5,231,633 | A | 7/1993  | Hluchyj et al.     | 6,081,522 | A  | 6/2000  | Hendel et al.         |
| 5,280,582 | A | 1/1994  | Yang et al.        | 6,088,356 | A  | 7/2000  | Hendel et al.         |
| 5,282,196 | A | 1/1994  | Clebowicz          | 6,094,434 | A  | 7/2000  | Kotzur et al.         |
| 5,287,477 | A | 2/1994  | Johnson et al.     | 6,101,552 | A  | 8/2000  | Chiang et al.         |
| 5,299,190 | A | 3/1994  | LaMaire et al.     | 6,104,696 | A  | 8/2000  | Kadambi et al.        |
| 5,299,195 | A | 3/1994  | Shah               | 6,104,700 | A  | 8/2000  | Haddock et al.        |
| 5,301,192 | A | 4/1994  | Henrion            | 6,104,969 | A  | 8/2000  | Beeks                 |
| 5,307,345 | A | 4/1994  | Lozowick et al.    | 6,108,306 | A  | 8/2000  | Kalkunte et al.       |
| 5,323,386 | A | 6/1994  | Wiher et al.       | 6,118,787 | A  | 9/2000  | Kalkunte et al.       |
| 5,365,512 | A | 11/1994 | Combs et al.       | 6,125,417 | A  | 9/2000  | Bailis et al.         |
| 5,377,189 | A | 12/1994 | Clark              | 6,128,666 | A  | 10/2000 | Muller et al.         |
| 5,390,173 | A | 2/1995  | Spinney et al.     | 6,144,668 | A  | 11/2000 | Bass et al.           |
| 5,392,279 | A | 2/1995  | Taniguchi          | 6,147,996 | A  | 11/2000 | Laor et al.           |
| 5,406,643 | A | 4/1995  | Burke et al.       | 6,151,301 | A  | 11/2000 | Holden                |
| 5,408,469 | A | 4/1995  | Opher et al.       | 6,151,497 | A  | 11/2000 | Yee et al.            |
| 5,430,442 | A | 7/1995  | Kaiser et al.      | 6,154,446 | A  | 11/2000 | Kadambi et al.        |
| 5,436,893 | A | 7/1995  | Barnett            | 6,157,643 | A  | 12/2000 | Ma                    |
| 5,461,615 | A | 10/1995 | Henrion            | 6,160,809 | A  | 12/2000 | Adiletta et al.       |
| 5,490,258 | A | 2/1996  | Fenner             | 6,160,812 | A  | 12/2000 | Bauman et al.         |
| 5,506,840 | A | 4/1996  | Pauwels et al.     | 6,172,990 | B1 | 1/2001  | Deb et al.            |
| 5,506,841 | A | 4/1996  | Sandquist          | 6,178,520 | B1 | 1/2001  | DeKoning et al.       |
| 5,521,923 | A | 5/1996  | Willmann et al.    | 6,181,699 | B1 | 1/2001  | Crinion et al.        |
| 5,539,733 | A | 7/1996  | Anderson et al.    | 6,185,208 | B1 | 2/2001  | Liao                  |
| 5,546,385 | A | 8/1996  | Caspi et al.       | 6,185,222 | B1 | 2/2001  | Hughes                |
| 5,550,816 | A | 8/1996  | Hardwick et al.    | 6,194,666 | B1 | 2/2001  | Hayshida et al.       |
| 5,563,948 | A | 10/1996 | Diehl et al.       | 6,195,335 | B1 | 2/2001  | Calvignac et al.      |
| 5,566,170 | A | 10/1996 | Bakke et al.       | 6,201,492 | B1 | 3/2001  | Amar et al.           |
| 5,598,410 | A | 1/1997  | Stone              | 6,222,845 | B1 | 4/2001  | Shue et al.           |
| 5,600,795 | A | 2/1997  | Du                 | 6,243,667 | B1 | 6/2001  | Kerr et al.           |
| 5,619,497 | A | 4/1997  | Gallagher et al.   | 6,249,528 | B1 | 6/2001  | Kothary               |
| 5,640,504 | A | 6/1997  | Johnson, Jr.       | 6,263,374 | B1 | 7/2001  | Olnowich et al.       |
| 5,646,878 | A | 7/1997  | Samra              | 6,272,144 | B1 | 8/2001  | Berenbaum et al.      |
| 5,649,089 | A | 7/1997  | Kilner             | 6,304,903 | B1 | 10/2001 | Ward                  |
| 5,663,952 | A | 9/1997  | Gentry, Jr.        | 6,320,859 | B1 | 11/2001 | Momirov               |
| 5,663,959 | A | 9/1997  | Nakagawa           | 6,333,929 | B1 | 12/2001 | Drottar et al.        |
| 5,666,353 | A | 9/1997  | Klausmeier et al.  | 6,335,932 | B2 | 1/2002  | Kadambi et al.        |
| 5,721,819 | A | 2/1998  | Galles et al.      | 6,335,935 | B2 | 1/2002  | Kadambi et al.        |
| 5,732,080 | A | 3/1998  | Ferguson et al.    | 6,343,072 | B1 | 1/2002  | Bechtolsheim et al.   |
| 5,734,826 | A | 3/1998  | Olnowich et al.    | 6,351,143 | B1 | 2/2002  | Guccione et al.       |
| 5,740,176 | A | 4/1998  | Gupta et al.       | 6,356,550 | B1 | 3/2002  | Williams              |
| 5,745,708 | A | 4/1998  | Weppler et al.     | 6,356,942 | B1 | 3/2002  | Bengtsson et al.      |
| 5,751,710 | A | 5/1998  | Crowther et al.    | 6,359,879 | B1 | 3/2002  | Carvey et al.         |
| 5,802,287 | A | 9/1998  | Rostoker et al.    | 6,363,077 | B1 | 3/2002  | Wong et al.           |
| 5,815,146 | A | 9/1998  | Youden et al.      | 6,366,557 | B1 | 4/2002  | Hunter                |
| 5,818,816 | A | 10/1998 | Chikazawa et al.   | 6,369,855 | B1 | 4/2002  | Chauvel et al.        |
| 5,835,496 | A | 11/1998 | Yeung et al.       | 6,370,579 | B1 | 4/2002  | Partridge             |
| 5,838,684 | A | 11/1998 | Wicki et al.       | 6,421,352 | B1 | 7/2002  | Manaka et al.         |
| 5,862,350 | A | 1/1999  | Coulson            | 6,424,658 | B1 | 7/2002  | Mathur                |
| 5,864,555 | A | 1/1999  | Mathur et al.      | 6,424,659 | B2 | 7/2002  | Viswanadham et al.    |
| 5,867,675 | A | 2/1999  | Lomelino et al.    | 6,427,185 | B1 | 7/2002  | Ryals et al.          |
| 5,870,538 | A | 2/1999  | Manning et al.     | 6,430,190 | B1 | 8/2002  | Essbaum et al.        |
| 5,872,769 | A | 2/1999  | Caldara et al.     | 6,442,067 | B1 | 8/2002  | Chawla et al.         |
| 5,872,783 | A | 2/1999  | Chin               | 6,457,175 | B1 | 9/2002  | Lerche                |
| 5,875,200 | A | 2/1999  | Glover et al.      | 6,459,705 | B1 | 10/2002 | Cheng                 |
| 5,896,380 | A | 4/1999  | Brown et al.       | 6,460,088 | B1 | 10/2002 | Merchant              |
| 5,907,566 | A | 5/1999  | Benson et al.      | 6,463,063 | B1 | 10/2002 | Bianchini, Jr. et al. |
|           |   |         |                    | 6,466,608 | B1 | 10/2002 | Hong et al.           |
|           |   |         |                    | 6,470,436 | B1 | 10/2002 | Croft et al.          |
|           |   |         |                    | 6,473,428 | B1 | 10/2002 | Nichols et al.        |
|           |   |         |                    | 6,473,433 | B1 | 10/2002 | Bianchini, Jr. et al. |



# US 8,811,390 B2

|           |    |         |                     |           |      |         |                    |         |
|-----------|----|---------|---------------------|-----------|------|---------|--------------------|---------|
| 6,477,174 | B1 | 11/2002 | Dooley et al.       | 6,842,903 | B1   | 1/2005  | Weschler           |         |
| 6,480,477 | B1 | 11/2002 | Treadaway et al.    | 6,854,117 | B1   | 2/2005  | Roberts            |         |
| 6,490,280 | B1 | 12/2002 | Leung               | 6,856,600 | B1   | 2/2005  | Russell et al.     |         |
| 6,493,347 | B2 | 12/2002 | Sindhu et al.       | 6,859,438 | B2   | 2/2005  | Haddock et al.     |         |
| 6,496,502 | B1 | 12/2002 | Fite, Jr. et al.    | 6,865,153 | B1   | 3/2005  | Hill et al.        |         |
| 6,505,281 | B1 | 1/2003  | Sherry              | 6,873,630 | B1   | 3/2005  | Muller et al.      |         |
| 6,510,138 | B1 | 1/2003  | Pannell             | 6,901,072 | B1 * | 5/2005  | Wong .....         | 370/389 |
| 6,522,656 | B1 | 2/2003  | Gridley             | 6,906,936 | B1   | 6/2005  | James et al.       |         |
| 6,532,229 | B1 | 3/2003  | Johnson et al.      | 6,912,637 | B1   | 6/2005  | Herbst             |         |
| 6,532,234 | B1 | 3/2003  | Yoshikawa et al.    | 6,920,154 | B1   | 7/2005  | Achler             |         |
| 6,535,504 | B1 | 3/2003  | Johnson et al.      | 6,925,516 | B2   | 8/2005  | Struhsaker et al.  |         |
| 6,549,519 | B1 | 4/2003  | Michels et al.      | 6,934,305 | B1   | 8/2005  | Duschatko et al.   |         |
| 6,553,370 | B1 | 4/2003  | Andreev et al.      | 6,937,606 | B2   | 8/2005  | Basso et al.       |         |
| 6,556,208 | B1 | 4/2003  | Congdon et al.      | 6,946,948 | B2   | 9/2005  | McCormack et al.   |         |
| 6,567,404 | B1 | 5/2003  | Wilford             | 6,952,419 | B1   | 10/2005 | Cassiday et al.    |         |
| 6,570,884 | B1 | 5/2003  | Connery et al.      | 6,957,258 | B2   | 10/2005 | Maher, III et al.  |         |
| 6,577,631 | B1 | 6/2003  | Keenan et al.       | 6,959,007 | B1   | 10/2005 | Vogel et al.       |         |
| 6,587,432 | B1 | 7/2003  | Putzolu et al.      | 6,961,347 | B1   | 11/2005 | Bunton et al.      |         |
| 6,591,302 | B2 | 7/2003  | Boucher et al.      | 6,965,615 | B1   | 11/2005 | Kerr et al.        |         |
| 6,601,186 | B1 | 7/2003  | Fox et al.          | 6,973,092 | B1   | 12/2005 | Zhou et al.        |         |
| 6,606,300 | B1 | 8/2003  | Blanc et al.        | 6,975,599 | B1   | 12/2005 | Johnson et al.     |         |
| 6,628,650 | B1 | 9/2003  | Saite et al.        | 6,978,309 | B1   | 12/2005 | Dorbolo            |         |
| 6,633,580 | B1 | 10/2003 | Torudbakken et al.  | 6,980,552 | B1   | 12/2005 | Belz et al.        |         |
| 6,636,483 | B1 | 10/2003 | Pannell             | 6,982,974 | B1   | 1/2006  | Saleh et al.       |         |
| 6,643,269 | B1 | 11/2003 | Fan et al.          | 6,990,102 | B1   | 1/2006  | Kaniz et al.       |         |
| 6,654,342 | B1 | 11/2003 | Dittia et al.       | 6,993,032 | B1   | 1/2006  | Dammann et al.     |         |
| 6,654,346 | B1 | 11/2003 | Mahalingaiah et al. | 6,996,663 | B1   | 2/2006  | Marsteiner         |         |
| 6,654,370 | B1 | 11/2003 | Quirke et al.       | 7,005,812 | B2   | 2/2006  | Mitchell           |         |
| 6,654,373 | B1 | 11/2003 | Maher, III et al.   | 7,009,968 | B2   | 3/2006  | Ambe et al.        |         |
| 6,658,002 | B1 | 12/2003 | Ross et al.         | 7,012,919 | B1   | 3/2006  | So et al.          |         |
| 6,661,791 | B1 | 12/2003 | Brown               | 7,012,919 | B1   | 3/2006  | So et al.          |         |
| 6,671,275 | B1 | 12/2003 | Wong et al.         | 7,050,505 | B2   | 5/2006  | Kaku               |         |
| 6,678,248 | B1 | 1/2004  | Haddock et al.      | 7,065,673 | B2   | 6/2006  | Subramaniam et al. |         |
| 6,681,332 | B1 | 1/2004  | Byrne et al.        | 7,082,133 | B1   | 7/2006  | Lor et al.         |         |
| 6,683,872 | B1 | 1/2004  | Saito               | 7,095,753 | B1   | 8/2006  | Milliken et al.    |         |
| 6,687,217 | B1 | 2/2004  | Chow et al.         | 7,103,041 | B1   | 9/2006  | Speiser et al.     |         |
| 6,687,247 | B1 | 2/2004  | Wilford et al.      | 7,106,692 | B1   | 9/2006  | Schulz             |         |
| 6,690,757 | B1 | 2/2004  | Bunton et al.       | 7,120,744 | B2   | 10/2006 | Klein              |         |
| 6,691,202 | B2 | 2/2004  | Vasquez et al.      | 7,126,948 | B2   | 10/2006 | Gooch et al.       |         |
| 6,696,917 | B1 | 2/2004  | Heitner et al.      | 7,126,956 | B2   | 10/2006 | Scholten           |         |
| 6,697,359 | B1 | 2/2004  | George              | 7,176,911 | B1   | 2/2007  | Kidono et al.      |         |
| 6,697,368 | B2 | 2/2004  | Chang et al.        | 7,185,141 | B1   | 2/2007  | James et al.       |         |
| 6,700,894 | B1 | 3/2004  | Shung               | 7,185,266 | B2   | 2/2007  | Blightman et al.   |         |
| 6,708,000 | B1 | 3/2004  | Nishi et al.        | 7,187,687 | B1   | 3/2007  | Davis et al.       |         |
| 6,721,229 | B1 | 4/2004  | Cole                | 7,188,237 | B2   | 3/2007  | Zhou et al.        |         |
| 6,721,268 | B1 | 4/2004  | Ohira et al.        | 7,190,696 | B1   | 3/2007  | Manur et al.       |         |
| 6,721,313 | B1 | 4/2004  | Van Duyne           | 7,191,277 | B2   | 3/2007  | Broyles            |         |
| 6,721,338 | B1 | 4/2004  | Sato                | 7,194,652 | B2   | 3/2007  | Zhou et al.        |         |
| 6,731,875 | B1 | 5/2004  | Kartalopoulos       | 7,203,194 | B2   | 4/2007  | Chang et al.       |         |
| 6,735,218 | B2 | 5/2004  | Chang et al.        | 7,206,283 | B2   | 4/2007  | Chang et al.       |         |
| 6,745,277 | B1 | 6/2004  | Lee et al.          | 7,212,536 | B2   | 5/2007  | Mackiewich et al.  |         |
| 6,747,971 | B1 | 6/2004  | Hughes et al.       | 7,218,637 | B1   | 5/2007  | Best et al.        |         |
| 6,751,224 | B1 | 6/2004  | Parruck et al.      | 7,219,293 | B2   | 5/2007  | Tsai et al.        |         |
| 6,754,881 | B2 | 6/2004  | Kuhlmann et al.     | 7,228,509 | B1   | 6/2007  | Dada et al.        |         |
| 6,765,866 | B1 | 7/2004  | Wyatt               | 7,236,490 | B2   | 6/2007  | Chang et al.       |         |
| 6,775,706 | B1 | 8/2004  | Fukumoto et al.     | 7,237,058 | B2   | 6/2007  | Srinivasan         |         |
| 6,778,546 | B1 | 8/2004  | Epps et al.         | 7,266,117 | B1   | 9/2007  | Davis              |         |
| 6,781,990 | B1 | 8/2004  | Puri et al.         | 7,272,611 | B1   | 9/2007  | Cuppett et al.     |         |
| 6,785,290 | B1 | 8/2004  | Fujisawa et al.     | 7,272,613 | B2   | 9/2007  | Sim et al.         |         |
| 6,788,697 | B1 | 9/2004  | Aweya et al.        | 7,277,425 | B1   | 10/2007 | Sikdar             |         |
| 6,792,484 | B1 | 9/2004  | Hook                | 7,283,547 | B1   | 10/2007 | Hook et al.        |         |
| 6,792,502 | B1 | 9/2004  | Pandya et al.       | 7,284,236 | B2   | 10/2007 | Zhou et al.        |         |
| 6,798,740 | B1 | 9/2004  | Senevirathne et al. | 7,286,534 | B2   | 10/2007 | Kloth              |         |
| 6,804,220 | B2 | 10/2004 | Odenwalder et al.   | 7,298,752 | B2   | 11/2007 | Moriwaki et al.    |         |
| 6,804,731 | B1 | 10/2004 | Chang et al.        | 7,324,509 | B2   | 1/2008  | Ni                 |         |
| 6,804,815 | B1 | 10/2004 | Kerr et al.         | 7,355,970 | B2   | 4/2008  | Lor                |         |
| 6,807,179 | B1 | 10/2004 | Kanuri et al.       | 7,356,030 | B2   | 4/2008  | Chang et al.       |         |
| 6,807,363 | B1 | 10/2004 | Abiko et al.        | 7,391,769 | B2   | 6/2008  | Rajkumar et al.    |         |
| 6,810,038 | B1 | 10/2004 | Isoyama et al.      | 7,414,979 | B1   | 8/2008  | Jarvis             |         |
| 6,810,046 | B2 | 10/2004 | Abbas et al.        | 7,428,693 | B2   | 9/2008  | Obuchi et al.      |         |
| 6,813,243 | B1 | 11/2004 | Epps et al.         | 7,468,975 | B1   | 12/2008 | Davis              |         |
| 6,813,266 | B1 | 11/2004 | Chiang et al.       | 7,512,127 | B2   | 3/2009  | Chang et al.       |         |
| 6,816,467 | B1 | 11/2004 | Muller et al.       | 7,543,077 | B1   | 6/2009  | Milliken et al.    |         |
| 6,831,923 | B1 | 12/2004 | Laor et al.         | 7,558,193 | B2   | 7/2009  | Bradbury et al.    |         |
| 6,831,932 | B1 | 12/2004 | Boyle et al.        | 7,561,590 | B1   | 7/2009  | Walsh              |         |
| 6,836,808 | B2 | 12/2004 | Bunce et al.        | 7,590,760 | B1   | 9/2009  | Banks              |         |
| 6,839,346 | B1 | 1/2005  | Kametani            | 7,596,139 | B2   | 9/2009  | Patel et al.       |         |
| 6,842,422 | B1 | 1/2005  | Bianchini, Jr.      | 7,606,968 | B2   | 10/2009 | Branscome et al.   |         |
|           |    |         |                     | 7,609,617 | B2   | 10/2009 | Appanna et al.     |         |
|           |    |         |                     | 7,613,991 | B1   | 11/2009 | Bain               |         |



# US 8,811,390 B2

|              |    |         |                     |              |    |         |                         |
|--------------|----|---------|---------------------|--------------|----|---------|-------------------------|
| 7,624,283    | B2 | 11/2009 | Bade et al.         | 2003/0152096 | A1 | 8/2003  | Chapman                 |
| 7,636,369    | B2 | 12/2009 | Wong                | 2003/0156586 | A1 | 8/2003  | Lee et al.              |
| 7,649,885    | B1 | 1/2010  | Davis et al.        | 2003/0159086 | A1 | 8/2003  | Arndt                   |
| 7,657,703    | B1 | 2/2010  | Singh               | 2003/0165160 | A1 | 9/2003  | Minami et al.           |
| 7,721,297    | B2 | 5/2010  | Ward                | 2003/0169470 | A1 | 9/2003  | Alagar et al.           |
| 7,738,450    | B1 | 6/2010  | Davis               | 2003/0174719 | A1 | 9/2003  | Sampath et al.          |
| 7,813,367    | B2 | 10/2010 | Wong                | 2003/0177221 | A1 | 9/2003  | Ould-Brahim et al.      |
| 7,817,659    | B2 | 10/2010 | Wong                | 2003/0198182 | A1 | 10/2003 | Pegrum et al.           |
| 7,821,972    | B1 | 10/2010 | Finn et al.         | 2003/0200343 | A1 | 10/2003 | Greenblat et al.        |
| 7,830,884    | B2 | 11/2010 | Davis               | 2003/0214956 | A1 | 11/2003 | Navada et al.           |
| 7,903,654    | B2 | 3/2011  | Bansal              | 2003/0215029 | A1 | 11/2003 | Limberg                 |
| 7,933,947    | B2 | 4/2011  | Fleischer et al.    | 2003/0223424 | A1 | 12/2003 | Anderson et al.         |
| 7,948,872    | B2 | 5/2011  | Patel et al.        | 2003/0223466 | A1 | 12/2003 | Noronha, Jr. et al.     |
| 7,953,922    | B2 | 5/2011  | Singh               | 2003/0227943 | A1 | 12/2003 | Hallman et al.          |
| 7,953,923    | B2 | 5/2011  | Singh               | 2004/0022263 | A1 | 2/2004  | Zhao et al.             |
| 7,978,614    | B2 | 7/2011  | Wong et al.         | 2004/0028060 | A1 | 2/2004  | Kang                    |
| 7,978,702    | B2 | 7/2011  | Chang et al.        | 2004/0054867 | A1 | 3/2004  | Stravers et al.         |
| 7,995,580    | B2 | 8/2011  | Patel et al.        | 2004/0062130 | A1 | 4/2004  | Chiang                  |
| 8,014,278    | B1 | 9/2011  | Subramanian et al.  | 2004/0062245 | A1 | 4/2004  | Sharp et al.            |
| 8,037,399    | B2 | 10/2011 | Wong et al.         | 2004/0062246 | A1 | 4/2004  | Boucher et al.          |
| 8,090,901    | B2 | 1/2012  | Lin et al.          | 2004/0083404 | A1 | 4/2004  | Subramaniam et al.      |
| 8,140,044    | B2 | 3/2012  | Villain et al.      | 2004/0088469 | A1 | 5/2004  | Levy                    |
| 8,149,839    | B1 | 4/2012  | Hsu et al.          | 2004/0128434 | A1 | 7/2004  | Khanna et al.           |
| 8,155,011    | B2 | 4/2012  | Wong et al.         | 2004/0141504 | A1 | 7/2004  | Blanc et al.            |
| 8,201,180    | B2 | 6/2012  | Briscoe et al.      | 2004/0179548 | A1 | 9/2004  | Chang et al.            |
| 8,238,255    | B2 | 8/2012  | Suresh et al.       | 2004/0190547 | A1 | 9/2004  | Gordy et al.            |
| 8,271,859    | B2 | 9/2012  | Wong et al.         | 2004/0205393 | A1 | 10/2004 | Kitamorn et al.         |
| 8,395,996    | B2 | 3/2013  | Wong et al.         | 2004/0208177 | A1 | 10/2004 | Ogawa                   |
| 8,448,162    | B2 | 5/2013  | Ramanathan et al.   | 2004/0208181 | A1 | 10/2004 | Clayton et al.          |
| 8,493,988    | B2 | 7/2013  | Wong et al.         | 2004/0223502 | A1 | 11/2004 | Wybenga et al.          |
| 8,509,236    | B2 | 8/2013  | Zhang et al.        | 2004/0235480 | A1 | 11/2004 | Rezaaifar et al.        |
| 2001/0001879 | A1 | 5/2001  | Kubik et al.        | 2004/0264380 | A1 | 12/2004 | Kalkunte et al.         |
| 2001/0007560 | A1 | 7/2001  | Masuda et al.       | 2005/0010630 | A1 | 1/2005  | Doering et al.          |
| 2001/0026551 | A1 | 10/2001 | Horlin              | 2005/0010849 | A1 | 1/2005  | Ryle et al.             |
| 2001/0048785 | A1 | 12/2001 | Steinberg           | 2005/0041684 | A1 | 2/2005  | Reynolds et al.         |
| 2001/0053150 | A1 | 12/2001 | Clear et al.        | 2005/0089049 | A1 | 4/2005  | Chang et al.            |
| 2002/0001307 | A1 | 1/2002  | Nguyen et al.       | 2005/0097432 | A1 | 5/2005  | Obuchi et al.           |
| 2002/0012585 | A1 | 1/2002  | Kalkunte et al.     | 2005/0120122 | A1 | 6/2005  | Farnham                 |
| 2002/0040417 | A1 | 4/2002  | Winograd et al.     | 2005/0132132 | A1 | 6/2005  | Rosenbluth et al.       |
| 2002/0048280 | A1 | 4/2002  | Lee et al.          | 2005/0132179 | A1 | 6/2005  | Glaum et al.            |
| 2002/0054594 | A1 | 5/2002  | Van Hoof et al.     | 2005/0138276 | A1 | 6/2005  | Navada et al.           |
| 2002/0054595 | A1 | 5/2002  | Ambe et al.         | 2005/0144369 | A1 | 6/2005  | Jaspers                 |
| 2002/0069294 | A1 | 6/2002  | Herkersdorf et al.  | 2005/0152324 | A1 | 7/2005  | Benveniste              |
| 2002/0073073 | A1 | 6/2002  | Cheng               | 2005/0152335 | A1 | 7/2005  | Lodha et al.            |
| 2002/0085499 | A1 | 7/2002  | Toyoyama et al.     | 2005/0169317 | A1 | 8/2005  | Pruecklmayer            |
| 2002/0087788 | A1 | 7/2002  | Morris              | 2005/0175018 | A1 | 8/2005  | Wong                    |
| 2002/0089937 | A1 | 7/2002  | Venkatachary et al. | 2005/0185577 | A1 | 8/2005  | Sakamoto et al.         |
| 2002/0089977 | A1 | 7/2002  | Chang et al.        | 2005/0185652 | A1 | 8/2005  | Iwamoto                 |
| 2002/0091844 | A1 | 7/2002  | Craft et al.        | 2005/0193316 | A1 | 9/2005  | Chen                    |
| 2002/0091884 | A1 | 7/2002  | Chang et al.        | 2005/0201387 | A1 | 9/2005  | Willis                  |
| 2002/0097713 | A1 | 7/2002  | Chang et al.        | 2005/0226236 | A1 | 10/2005 | Klink                   |
| 2002/0105966 | A1 | 8/2002  | Patel et al.        | 2005/0246508 | A1 | 11/2005 | Shaw                    |
| 2002/0126672 | A1 | 9/2002  | Chow et al.         | 2005/0249124 | A1 | 11/2005 | Elie-Dit-Cosaque et al. |
| 2002/0131437 | A1 | 9/2002  | Tagore-Brage        | 2006/0031610 | A1 | 2/2006  | Liav et al.             |
| 2002/0141403 | A1 | 10/2002 | Akahane et al.      | 2006/0034452 | A1 | 2/2006  | Tomomura                |
| 2002/0146013 | A1 | 10/2002 | Karlsson et al.     | 2006/0050690 | A1 | 3/2006  | Epps et al.             |
| 2002/0161967 | A1 | 10/2002 | Kirihata et al.     | 2006/0077891 | A1 | 4/2006  | Smith et al.            |
| 2002/0169786 | A1 | 11/2002 | Richek              | 2006/0092829 | A1 | 5/2006  | Brolin et al.           |
| 2002/0191605 | A1 | 12/2002 | Van Lunteren et al. | 2006/0092929 | A1 | 5/2006  | Chun                    |
| 2003/0009466 | A1 | 1/2003  | Ta et al.           | 2006/0114876 | A1 | 6/2006  | Kalkunte                |
| 2003/0012198 | A1 | 1/2003  | Kaganoi et al.      | 2006/0146374 | A1 | 7/2006  | Ng et al.               |
| 2003/0033435 | A1 | 2/2003  | Hanner              | 2006/0165089 | A1 | 7/2006  | Klink                   |
| 2003/0043800 | A1 | 3/2003  | Sonksen et al.      | 2006/0209685 | A1 | 9/2006  | Rahman et al.           |
| 2003/0043848 | A1 | 3/2003  | Sonksen             | 2006/0221841 | A1 | 10/2006 | Lee et al.              |
| 2003/0048785 | A1 | 3/2003  | Calvignac et al.    | 2006/0268680 | A1 | 11/2006 | Roberts et al.          |
| 2003/0061459 | A1 | 3/2003  | Aboulenein et al.   | 2006/0274749 | A1 | 12/2006 | Beier                   |
| 2003/0074657 | A1 | 4/2003  | Bramley, Jr.        | 2007/0038798 | A1 | 2/2007  | Bouchard et al.         |
| 2003/0081608 | A1 | 5/2003  | Barri et al.        | 2007/0088974 | A1 | 4/2007  | Chandwani et al.        |
| 2003/0095548 | A1 | 5/2003  | Yamano              | 2007/0127464 | A1 | 6/2007  | Jain et al.             |
| 2003/0103499 | A1 | 6/2003  | Davis et al.        | 2007/0179909 | A1 | 8/2007  | Channasagara            |
| 2003/0103500 | A1 | 6/2003  | Menon et al.        | 2007/0208876 | A1 | 9/2007  | Davis                   |
| 2003/0108052 | A1 | 6/2003  | Inoue et al.        | 2007/0258475 | A1 | 11/2007 | Chinn et al.            |
| 2003/0110180 | A1 | 6/2003  | Calvignac et al.    | 2007/0288690 | A1 | 12/2007 | Wang et al.             |
| 2003/0115403 | A1 | 6/2003  | Bouchard et al.     | 2008/0002707 | A1 | 1/2008  | Davis                   |
| 2003/0120861 | A1 | 6/2003  | Calle et al.        | 2008/0025309 | A1 | 1/2008  | Swallow                 |
| 2003/0128668 | A1 | 7/2003  | Yavatkar et al.     | 2008/0031263 | A1 | 2/2008  | Ervin et al.            |
| 2003/0137978 | A1 | 7/2003  | Kanetake            | 2008/0037544 | A1 | 2/2008  | Yano et al.             |
| 2003/0152084 | A1 | 8/2003  | Lee et al.          | 2008/0049742 | A1 | 2/2008  | Bansal et al.           |
|              |    |         |                     | 2008/0069125 | A1 | 3/2008  | Reed et al.             |



|              |    |         |                    |
|--------------|----|---------|--------------------|
| 2008/0092020 | A1 | 4/2008  | Hasenplaugh et al. |
| 2008/0095169 | A1 | 4/2008  | Chandra et al.     |
| 2008/0117075 | A1 | 5/2008  | Seddigh et al.     |
| 2008/0126652 | A1 | 5/2008  | Vembu et al.       |
| 2008/0159309 | A1 | 7/2008  | Sultan et al.      |
| 2008/0181103 | A1 | 7/2008  | Davies             |
| 2008/0205407 | A1 | 8/2008  | Chang et al.       |
| 2008/0307288 | A1 | 12/2008 | Ziesler et al.     |
| 2009/0175178 | A1 | 7/2009  | Yoon et al.        |
| 2009/0207838 | A1 | 8/2009  | Milliken et al.    |
| 2009/0279423 | A1 | 11/2009 | Suresh et al.      |
| 2009/0279440 | A1 | 11/2009 | Wong et al.        |
| 2009/0279441 | A1 | 11/2009 | Wong et al.        |
| 2009/0279541 | A1 | 11/2009 | Wong et al.        |
| 2009/0279542 | A1 | 11/2009 | Wong et al.        |
| 2009/0279546 | A1 | 11/2009 | Davis              |
| 2009/0279548 | A1 | 11/2009 | Davis et al.       |
| 2009/0279549 | A1 | 11/2009 | Ramanathan et al.  |
| 2009/0279558 | A1 | 11/2009 | Davis et al.       |
| 2009/0279559 | A1 | 11/2009 | Wong et al.        |
| 2009/0279561 | A1 | 11/2009 | Chang et al.       |
| 2009/0282148 | A1 | 11/2009 | Wong et al.        |
| 2009/0282322 | A1 | 11/2009 | Wong et al.        |
| 2009/0287952 | A1 | 11/2009 | Patel et al.       |
| 2009/0290499 | A1 | 11/2009 | Patel et al.       |
| 2010/0034215 | A1 | 2/2010  | Patel et al.       |
| 2010/0046521 | A1 | 2/2010  | Wong               |
| 2010/0100671 | A1 | 4/2010  | Singh              |
| 2010/0135313 | A1 | 6/2010  | Davis              |
| 2010/0161894 | A1 | 6/2010  | Singh              |
| 2010/0246588 | A1 | 9/2010  | Davis              |
| 2010/0293327 | A1 | 11/2010 | Lin et al.         |
| 2011/0002340 | A1 | 1/2011  | Davis              |
| 2011/0044340 | A1 | 2/2011  | Bansal et al.      |
| 2011/0069711 | A1 | 3/2011  | Jha et al.         |
| 2011/0110237 | A1 | 5/2011  | Wong et al.        |
| 2011/0173386 | A1 | 7/2011  | Milliken et al.    |
| 2012/0023309 | A1 | 1/2012  | Abraham et al.     |
| 2012/0026868 | A1 | 2/2012  | Chang et al.       |
| 2012/0163389 | A1 | 6/2012  | Zhang et al.       |
| 2012/0236722 | A1 | 9/2012  | Patel et al.       |
| 2012/0294312 | A1 | 11/2012 | Edward et al.      |
| 2013/0034098 | A1 | 2/2013  | Davis              |

FOREIGN PATENT DOCUMENTS

|    |             |    |         |
|----|-------------|----|---------|
| JP | 2004-537871 | A  | 12/2004 |
| WO | WO 01/84728 | A1 | 11/2001 |
| WO | WO 02/41544 | A3 | 5/2007  |

OTHER PUBLICATIONS

Non-Final Office Action for U.S. Appl. No. 11/831,950, mailed Aug. 26, 2011, 45 pages.  
 Final Office Action for U.S. Appl. No. 11/953,742, mailed on Oct. 26, 2011, 19 pages.  
 Non-Final Office Action for U.S. Appl. No. 11/668,322, mailed on Aug. 30, 2011 17 pages.  
 Non-Final Office Action for U.S. Appl. No. 11/745,008, mailed on Sep. 14, 2011, 26 pages.  
 Notice of Allowance for U.S. Appl. No. 12/795,492, mailed on Nov. 14, 2011, 10 pages.  
 Final Office Action for U.S. Appl. No. 12/198,710, mailed on Oct. 19, 2011, 58 pages.  
 Final Office Action for U.S. Appl. No. 12/070,893, mailed on Sep. 21, 2011, 12 pages.  
 Notice of Allowance for U.S. Appl. No. 12/466,277, mailed on Nov. 2, 2011, 47 pages.  
 Non-Final Office Action for U.S. Appl. No. 11/953,742, mailed on Mar. 30, 2011, 23 pages.  
 Non-Final Office Action for U.S. Appl. No. 11/953,751, mailed on Mar. 29, 2011, 29 pages.  
 Non-Final Office Action for U.S. Appl. No. 11/779,778, mailed on Feb. 2, 2011, 63 pages.  
 Final Office Action for U.S. Appl. No. 11/668,322, mailed on Feb. 1, 2011, 17 pages.

Non-Final Office Action for U.S. Appl. No. 12/795,492, mailed on Mar. 17, 2011, 51 pages.  
 Non-Final Office Action for U.S. Appl. No. 12/198,710, mailed on Mar. 24, 2011, 40 pages.  
 Notice of Allowance for U.S. Appl. No. 12/639,749, mailed on Feb. 11, 2011, 51 pages.  
 Notice of Allowance for U.S. Appl. No. 12/639,762, mailed on Mar. 4, 2011, 7 pages.  
 Notice of Allowance for U.S. Appl. No. 12/400,591, mailed on Mar. 23, 2011, 11 pages.  
 Notice of Allowance for U.S. Appl. No. 12/372,390, mailed Mar. 9, 2011, 8 pages.  
 Non-Final Office Action for U.S. Appl. No. 12/070,893, mailed on Mar. 18, 2011, 7 pages.  
 Advisory Action for U.S. Appl. No. 11/831,950, mailed on Mar. 4, 2010, 4 pages.  
 Non-Final Office Action for U.S. Appl. No. 11/779,714, mailed on Mar. 31, 2010, 26 pages.  
 Non-Final Office Action for U.S. Appl. No. 10/810,208, mailed on Feb. 5, 2010, 13 pages.  
 Final Office Action for U.S. Appl. No. 11/668,322, mailed on Feb. 24, 2010, 33 pages.  
 Final Office Action for U.S. Appl. No. 10/140,751, mailed on Mar. 25, 2010, 29 pages.  
 Advisory Action for U.S. Appl. No. 11/745,008, mailed on Apr. 21, 2010, 8 pages.  
 Notice of Allowance for U.S. Appl. No. 11/621,038, mailed on Apr. 28, 2010, 15 pages.  
 Non-Final Office Action for U.S. Appl. No. 12/400,594, mailed on May 14, 2010, 53 pages.  
 Non-Final Office Action for U.S. Appl. No. 12/372,390, mailed on Apr. 22, 2010, 46 pages.  
 Davis, et al., "Pipeline Method and System for Switching Packets," U.S. Appl. No. 13/398,725, filed Feb. 16, 2012, 51 pages.  
 Hsu, et al., "Selection of Trunk Ports and Paths Using Rotation," U.S. Appl. No. 13/407,397, filed Feb. 28, 2012, 30 pages.  
 Final Office Action for U.S. Appl. No. 11/831,950, mailed Feb. 28, 2012, 20 pages.  
 Notice of Allowance for U.S. Appl. No. 11/831,950, mailed May 16, 2012, 11 pages.  
 Notice of Allowance for U.S. Appl. No. 11/953,751, mailed Dec. 7, 2011, 12 pages.  
 Supplemental Notice of Allowance for U.S. Appl. No. 11/953,751, mailed Dec. 27, 2011, 6 pages.  
 Notice of Allowance for U.S. Appl. No. 11/668,322, mailed Feb. 10, 2012, 20 pages.  
 Notice of Allowance for U.S. Appl. No. 12/198,697, mailed Nov. 28, 2011, 12 pages.  
 Notice of Allowance for U.S. Appl. No. 12/198,697, mailed Jan. 5, 2012, 4 pages.  
 Non-Final Office Action for U.S. Appl. No. 13/083,481, mailed on Dec. 1, 2011, 7 pages.  
 Final Office Action for U.S. Appl. No. 11/646,845, mailed on Jun. 9, 2011, 22 pages.  
 Notice of Allowance for U.S. Appl. No. 11/953,743, mailed on Apr. 28, 2011, 19 pages.  
 Final Office Action for U.S. Appl. No. 10/140,751, mailed on Jun. 28, 2011, 23 pages.  
 Non-Final Office Action for U.S. Appl. No. 12/702,031, mailed on Apr. 29, 2011, 5 pages.  
 Non-Final Office Action for U.S. Appl. No. 12/198,697, mailed on May 20, 2011, 43 pages.  
 Final Office Action for U.S. Appl. No. 11/953,742, mailed on Jun. 14, 2010, 21 pages.  
 Final Office Action for U.S. Appl. No. 11/953,743, mailed on Jul. 15, 2010, 21 pages.  
 Non-Final Office Action for U.S. Appl. No. 11/953,745, mailed on Jun. 14, 2010, 19 pages.  
 Final Office Action for U.S. Appl. No. 11/953,751, mailed on Jun. 25, 2010, 24 pages.  
 Notice of Allowance for U.S. Appl. No. 10/810,208, mailed on Jul. 15, 2010, 15 pages.



(56)

**References Cited**

## OTHER PUBLICATIONS

- Non-Final Office Action for U.S. Appl. No. 11/668,322, mailed on Jun. 22, 2010, 16 pages.
- Notice of Allowance for U.S. Appl. No. 11/854,486, mailed on Jul. 13, 2010, 12 pages.
- Final Office Action for U.S. Appl. No. 12/198,697, mailed on Aug. 2, 2010, 55 pages.
- Non-Final Office Action for U.S. Appl. No. 12/070,893, mailed on Jun. 10, 2010, 44 pages.
- Advisory Action for U.S. Appl. No. 11/615,769, mailed on May 25, 2010, 3 pages.
- Notice of Allowance for U.S. Appl. No. 11/615,769, mailed on Jul. 12, 2010, 14 pages.
- Notice of Allowance for U.S. Appl. No. 11/779,778, mailed on Jul. 28, 2011, 9 pages.
- Final Office Action for U.S. Appl. No. 12/795,492, mailed on Jul. 20, 2011, 11 pages.
- Requirement for Restriction/Election for U.S. Appl. No. 12/466,277, mailed on Aug. 9, 2011, 6 pages.
- Non-Final Office Action for U.S. Appl. No. 12/900,279 mailed on Apr. 11, 2013, 7 pages.
- Gupta et al., "Packet Classification on Multiple Fields," SIGCOMM '99, Aug. 1999, pp. 147-160, ACM, Cambridge, MA.
- Final Office Action for U.S. Appl. No. 11/745,008, mailed on Jun. 28, 2012.
- Final Office Action for U.S. Appl. No. 11/646,845, mailed on Jul. 5, 2012, 21 pages.
- Non-Final Office Action for U.S. Appl. No. 12/900,279 mailed Aug. 30, 2012, 9 pages.
- Non-Final Office Action for U.S. Appl. No. 12/198,710 mailed Sep. 13, 2012, 18 pages.
- Non-Final Office Action for U.S. Appl. No. 13/083,481 mailed Oct. 4, 2012, 9 pages.
- Non-Final Office Action for U.S. Appl. No. 12/880,518 mailed Oct. 30, 2012, 7 pages.
- Non-Final Office Action for U.S. Appl. No. 13/152,715 mailed on Nov. 13, 2012, 6 pages.
- Notice of Allowance for U.S. Appl. No. 11/953,742 mailed on Nov. 13, 2012, 7 pages.
- Non-Final Office Action for U.S. Appl. No. 13/398,725 mailed on Nov. 28, 2012, 10 pages.
- Final Office Action for U.S. Appl. No. 12/900,279 mailed on Dec. 5, 2012, 8 pages.
- Non-Final Office Action for U.S. Appl. No. 12/624,300 mailed on Dec. 31, 2012, 13 pages.
- Notice of Allowance for U.S. Appl. No. 11/646,845 mailed on Jan. 8, 2013, 8 pages.
- Non-Final Office Action for U.S. Appl. No. 12/684,022 mailed Jul. 30, 2012, 18 pages.
- Notice of Allowance for U.S. Appl. No. 11/779,714, mailed on Jun. 18, 2012, 9 pages.
- Final Office Action for U.S. Appl. No. 12/198,710 mailed on Mar. 21, 2013, 17 pages.
- Non-Final Office Action for U.S. Appl. No. 13/083,481 mailed on Mar. 1, 2013, 14 pages.
- Non-Final Office Action for U.S. Appl. No. 11/745,008 mailed on Mar. 7, 2013, 18 pages.
- Non-Final Office Action for U.S. Appl. No. 13/548,116 mailed on Apr. 15, 2013, 8 pages.
- U.S. Appl. No. Appl. 12/639,762, filed Dec. 16, 2009, Singh.
- U.S. Appl. No. Appl. 12/639,749, filed Dec. 16, 2009, Singh.
- U.S. Appl. No. Appl. 12/624,300, filed Nov. 23, 2009, Davis et al.
- U.S. Appl. No. Appl. 12/608,985, filed Oct. 29, 2009, Wong.
- U.S. Appl. No. Appl. 12/505,390, filed Jul. 17, 2009, Patel et al.
- U.S. Appl. No. Appl. 12/417,913, filed Apr. 3, 2009, Patel et al.
- U.S. Appl. No. Appl. 12/198,697, filed Aug. 26, 2008, Hsu et al.
- U.S. Appl. No. Appl. 11/828,246, filed Jul. 25, 2007, Davis.
- U.S. Appl. No. Appl. 11/724,965, filed Mar. 15, 2007, Chang et al.
- U.S. Appl. No. Appl. 11/586,991, filed Oct. 25, 2006, Ramanathan et al.
- U.S. Appl. No. Appl. 10/832,086, filed Apr. 26, 2004, Wong.
- U.S. Appl. No. Appl. 10/141,223, filed May 7, 2002, Veerabadran et al.
- U.S. Appl. No. Appl. 10/140,753, filed May 6, 2002, David et al.
- U.S. Appl. No. Appl. 10/140,751, filed May 6, 2002, Davis.
- 10 Gigabit Ethernet—Technology Overview White Paper, Sep. 2001, 16 pages.
- 10 Gigabit Ethernet Alliance, Interconnection with Wide Area Networks, Version 1.0, Mar. 2002, 6 pages.
- "IEEE Standard for Information Technology—Telecommunications and Information Exchange Between Systems—Local and Metropolitan Area Networks—Common Specifications—Part 3: Media Access Control (MAC) Bridges," ANSI/IEEE Standard 802.1D, 1998 Edition, 373 pages.
- Belhadj, et al., "Feasibility of a 100GE MAC," PowerPoint Presentation, IEEE Meeting Nov. 2006, Nov. 13-15, 2006, 18 pages.
- Braun, et al., "Fast incremental CRC updates for IP over ATM networks," IEEE Workshop on High Performance Switching and Routing, 2001, 6 pages.
- Degermark, M., et al., "Small Forwarding Tables for Fast Routing Lookups," ACM Computer Communications Review 27(4):3-14, Oct. 1997.
- Foundry Networks, "BigIron Architecture Technical Brief," Dec. 1998—Version 1.03, 14 pages.
- Foundry Networks, "BigIron Architecture Technical Brief," Jul. 2001—Version 2.02, 16 pages.
- Foundry Networks, "BigIron Architecture Technical Brief," May 1999—Version 2.0, 15 pages.
- Foundry Networks, "BigIron Architecture Technical Brief," May, 1999—Version 2.01, 15 pages.
- Foundry Networks, "BigIron Architecture Technical Brief," Oct. 1998—Version 1.0, 15 pages.
- Foundry Networks, "BigIron Architecture Technical Brief," Oct. 1998—Version 1.02, 15 pages.
- Foundry Networks, "Foundry Networks, Next Generation Terabit System Architecture —The High Performance Revolution for 10 Gigabit Networks," Nov. 17, 2003, 27 pages.
- Foundry Networks, "JetCore™ Based Chassis Systems —An Architecture Brief on NetIron, BigIron, and FastIron Systems," Jan. 17, 2003, 27 pages.
- Gigabit Ethernet Alliance—"Accelerating the Standard for Speed," Copyright 1998, 19 pages.
- Kichorowsky, R., et al., "Mindspeed.TM. Switch Fabric Offers the Most Comprehensive Solution for Multi-Protocol Networking Equipment," Apr. 30, 2001, 3 pages.
- Matsumoto, C., et al., "Switch Fabrics Touted at Interconnects Conference," Aug. 21, 2000, URL= <http://www.eetimes.com/story/OEG2000821S0011>, accessed Aug. 12, 2002, 2 pages.
- McAuley, et al., "Fast Routing Table Lookup Using CAMs," Proceedings of INFOCOM, Mar.-Apr. 1993, pp. 1382-1391.
- Mier Communications, Inc., "Lab Testing Summary Report—Product Category: Layer-3 Switches, Vendor Tested:, Product Tested: Foundry Networks, BigIron 4000," Report No. 231198, Oct. 1998, 6 pages.
- Mier Communications, Inc., "Lab Testing Summary Report—Product Category: Gigabit Backbone Switches, Vendor Tested: Foundry Networks, Product Tested: BigIron 4000," Report No. 210998, Sep. 1998, 6 pages.
- Mindspeed—A Conexant Business, "17x17 3.2 Gbps Crosspoint Switch with Input Equalization—M21110," Feb. 1, 2001, 2 pages.
- Mindspeed—A Conexant Business, "Switch Fabric Chipset—CX27300 iScale.TM.," Apr. 30, 2001, 2 pages.
- Newton, Newton's Telecom Dictionary, CMP Books, Mar. 2004, 20th Ed., p. 617.
- Satran et al., "Out of Order Incremental CRC Computation," IEEE Transactions on Computers, vol. 54, Issue 9 Sep. 2005), 11 pages.
- Spurgeon, "Ethernet, The Definitive Guide," O'Reilly & Associates, Inc., Sebastapol, CA, Feb. 2000.
- The Tolly Group, "Foundry Networks, Inc.—BigIron 4000, Layer 2 & Layer 3 Interoperability Evaluation," No. 199133, Oct. 1999, 4 pages.



(56)

**References Cited**

## OTHER PUBLICATIONS

The Tolly Group, "Foundry Networks, Inc.—BigIron 8000 Gigabit Ethernet Switching Router, Layer 2 & Layer 3 Performance Evaluation," No. 199111, May, 1999, 4 pages.

International Search Report for Application No. PCT/US2001/043113, mailed Dec. 13, 2002, 2 pages.

Written Opinion of the International Searching Authority for Application No. PCT/US2001/043113, mailed May 1, 2003, 6 pages.

International Preliminary Examination Report for Application No. PCT/US2001/043113, mailed Nov. 6, 2003, 6 pages.

International Search Report for Application No. PCT/US03/08719, mailed Jun. 19, 2003, 1 page.

Non-Final Office Action for U.S. Appl. No. 09/855,024, mailed Jun. 4, 2002, 10 pages.

Final Office Action for U.S. Appl. No. 09/855,024, mailed Jan. 15, 2003, 20 pages.

Advisory Action for U.S. Appl. No. 09/855,024, mailed May 2, 2003.

Notice of Allowance for U.S. Appl. No. 09/855,024, mailed Nov. 3, 2003.

Notice of Allowance for U.S. Appl. No. 09/855,024, mailed Dec. 15, 2003, 6 pages.

Non-Final Office Action for U.S. Appl. No. 10/810,301, mailed Mar. 17, 2005, 11 pages.

Non-Final Office Action for U.S. Appl. No. 10/810,301, mailed Feb. 16, 2006, 12 pages.

Notice of Allowance for U.S. Appl. No. 10/810,301, mailed Jul. 28, 2006, 5 pages.

Notice of Allowance for U.S. Appl. No. 10/810,301, mailed Feb. 6, 2007, 9 pages.

Non-Final Office Action for U.S. Appl. No. 09/855,025, mailed Nov. 23, 2004, 17 pages.

Non-Final Office Action for U.S. Appl. No. 09/855,031, mailed May 22, 2002.

Non-Final Office Action for U.S. Appl. No. 09/855,031, mailed Dec. 10, 2002.

Final Office Action for U.S. Appl. No. 09/855,031, mailed Jul. 30, 2003.

Notice of Allowance for U.S. Appl. No. 09/855,031, mailed Nov. 4, 2003.

Non-Final Office Action for U.S. Appl. No. 10/736,680, mailed Feb. 16, 2006, 18 pages.

Final Office Action for U.S. Appl. No. 10/736,680, mailed Aug. 3, 2006, 10 pages.

Notice of Allowance for U.S. Appl. No. 10/736,680, mailed Feb. 22, 2007, 12 pages.

Non-Final Office Action for U.S. Appl. No. 10/210,041, mailed Sep. 10, 2003, 12 pages.

Final Office Action for U.S. Appl. No. 10/210,041, mailed Jan. 7, 2004, 14 pages.

Non-Final Office Action for U.S. Appl. No. 10/210,041, mailed Mar. 11, 2004, 12 pages.

Final Office Action for U.S. Appl. No. 10/210,041, mailed Jul. 7, 2004, 13 pages.

Non-Final Office Action for U.S. Appl. No. 10/210,041, mailed Feb. 9, 2005, 7 pages.

Final Office Action for U.S. Appl. No. 10/210,041, mailed Aug. 24, 2005, 7 pages.

Advisory Action for U.S. Appl. No. 10/210,041, mailed Dec. 13, 2005, 4 pages.

Non-Final Office Action for U.S. Appl. No. 10/210,108, mailed Jun. 12, 2003, 6 pages.

Notice of Allowance for U.S. Appl. No. 10/210,108, mailed Oct. 7, 2003.

Requirement for Restriction/Election for U.S. Appl. No. 10/438,545, mailed Oct. 31, 2003.

Non-Final Office Action for U.S. Appl. No. 10/438,545, mailed Dec. 12, 2003, 7 pages.

Notice of Allowance for U.S. Appl. No. 10/438,545, mailed Jun. 15, 2004, 4 pages.

Non-Final Office Action for U.S. Appl. No. 10/832,086, mailed Sep. 19, 2007, 12 pages.

Final Office Action for U.S. Appl. No. 10/832,086, mailed May 1, 2008, 31 pages.

Advisory Action for U.S. Appl. No. 10/832,086, mailed Jul. 21, 2008, 4 pages.

Non-Final Office Action for U.S. Appl. No. 10/832,086, mailed Sep. 18, 2008, 18 pages.

Non-Final Office Action for U.S. Appl. No. 10/832,086, mailed Apr. 1, 2009, 17 pages.

Final Office Action for U.S. Appl. No. 10/832,086, mailed Sep. 29, 2009, 26 pages.

Non-Final Office Action for U.S. Appl. No. 11/586,991, mailed Oct. 2, 2008, 23 pages.

Non-Final Office Action for U.S. Appl. No. 11/831,950, mailed Aug. 18, 2009, 49 pages.

Final Office Action for U.S. Appl. No. 11/831,950, mailed Jan. 6, 2010, 21 pages.

Non-Final Office Action for U.S. Appl. No. 11/953,742, mailed Nov. 19, 2009, 51 pages.

Non-Final Office Action for U.S. Appl. No. 11/953,743, mailed Nov. 23, 2009, 47 pages.

Non-Final Office Action for U.S. Appl. No. 11/953,745, mailed Nov. 24, 2009, 48 pages.

Non-Final Office Action for U.S. Appl. No. 11/953,751, mailed Nov. 16, 2009, 55 pages.

Non-Final Office Action for U.S. Appl. No. 11/779,714, mailed Sep. 1, 2009, 58 pages.

Non-Final Office Action for U.S. Appl. No. 10/810,208, mailed Jul. 16, 2007, 24 pages.

Non-Final Office Action for U.S. Appl. No. 10/810,208, mailed Dec. 18, 2007, 40 pages.

Final Office Action for U.S. Appl. No. 10/810,208, mailed Jun. 11, 2008, 34 pages.

Advisory Action for U.S. Appl. No. 10/810,208, mailed Aug. 27, 2008, 4 pages.

Non-Final Office Action for U.S. Appl. No. 10/810,208, mailed Feb. 13, 2009, 17 pages.

Non-Final Office Action for U.S. Appl. No. 10/810,208, mailed Aug. 24, 2009, 38 pages.

Requirement for Restriction/Election for U.S. Appl. No. 10/140,752, mailed May 18, 2006, 8 pages.

Non-Final Office Action for U.S. Appl. No. 10/140,752, mailed Dec. 14, 2006, 17 pages.

Non-Final Office Action for U.S. Appl. No. 10/140,752, mailed Apr. 23, 2007, 6 pages.

Non-Final Office Action for U.S. Appl. No. 10/140,752, mailed Jan. 24, 2008, 8 pages.

Notice of Allowance of U.S. Appl. No. 10/140,752, mailed Jul. 24, 2008, 14 pages.

Notice of Allowance of U.S. Appl. No. 10/140,752, mailed Sep. 10, 2008, 4 pages.

Non-Final Office Action for U.S. Appl. No. 11/668,322, mailed Mar. 23, 2009, 19 pages.

Requirement for Restriction/Election for U.S. Appl. No. 11/668,322, mailed Oct. 29, 2009, 6 pages.

Non-Final Office Action for U.S. Appl. No. 11/854,486, mailed Jul. 20, 2009, 29 pages.

Non-Final Office Action for U.S. Appl. No. 11/854,486, mailed Jan. 12, 2010, 23 pages.

Non-Final Office Action for U.S. Appl. No. 10/139,912, mailed Jan. 25, 2006, 14 pages.

Final Office Action for U.S. Appl. No. 10/139,912, mailed Aug. 11, 2006, 26 pages.

Non-Final Office Action for U.S. Appl. No. 10/139,912, mailed Apr. 20, 2007, 20 pages.

Final Office Action for U.S. Appl. No. 10/139,912, mailed Nov. 28, 2007, 20 pages.

Non-Final Office Action for U.S. Appl. No. 10/139,912, mailed Aug. 1, 2008, 21 pages.

Notice of Allowance for U.S. Appl. No. 10/139,912, mailed Feb. 5, 2009, 8 pages.



(56)

**References Cited**

## OTHER PUBLICATIONS

Notice of Allowance for U.S. Appl. No. 10/139,912, mailed Jun. 8, 2009, 8 pages.

Notice of Allowance for U.S. Appl. No. 10/139,912, mailed on Oct. 19, 2009, 17 pages.

Supplemental Notice of Allowance for U.S. Appl. No. 10/139,912, mailed on Nov. 23, 2009, 4 pages.

Requirement for Restriction/Election for U.S. Appl. No. 10/140,751, mailed Apr. 27, 2006, 5 pages.

Non-Final Office Action for U.S. Appl. No. 10/140,751, mailed Aug. 10, 2006, 15 pages.

Final Office Action for U.S. Appl. No. 10/140,751, mailed Apr. 10, 2007, 16 pages.

Non-Final Office Action for U.S. Appl. No. 10/140,751, mailed Oct. 30, 2007, 14 pages.

Final Office Action for U.S. Appl. No. 10/140,751, mailed May 28, 2008, 19 pages.

Non-Final Office Action for U.S. Appl. No. 10/140,751, mailed Sep. 17, 2008, 15 pages.

Final Office Action for U.S. Appl. No. 10/140,751, mailed Mar. 17, 2009, 17 pages.

Advisory Action for U.S. Appl. No. 10/140,751, mailed Jun. 1, 2009, 3 pages.

Non-Final Office Action for U.S. Appl. No. 10/140,751, mailed Sep. 28, 2009, 34 pages.

Non-Final Office Action for U.S. Appl. No. 11/745,008, mailed May 14, 2009, 27 pages.

Final Office Action for U.S. Appl. No. 11/745,008, mailed Dec. 30, 2009, 27 pages.

Non-Final Office Action for U.S. Appl. No. 10/141,223, mailed Feb. 23, 2006, 25 pages.

Non-Final Office Action for U.S. Appl. No. 10/141,223, mailed Feb. 13, 2007, 29 pages.

Final Office Action for U.S. Appl. No. 10/141,223, mailed Aug. 21, 2007, 25 pages.

Non-Final Office Action for U.S. Appl. No. 10/141,223, mailed Dec. 28, 2007, 13 pages.

Non-Final Office Action for U.S. Appl. No. 10/141,223, mailed Sep. 3, 2008, 22 pages.

Non-Final Office Action for U.S. Appl. No. 10/139,831, mailed Oct. 17, 2005, 7 pages.

Notice of Allowance for U.S. Appl. No. 10/139,831, mailed Feb. 9, 2006, 7 pages.

Non-Final Office Action for U.S. Appl. No. 10/139,831, mailed Jun. 27, 2006, 9 pages.

Final Office Action for U.S. Appl. No. 10/139,831, mailed Nov. 28, 2006, 17 pages.

Notice of Allowance for U.S. Appl. No. 10/139,831, mailed Jun. 14, 2007, 26 pages.

Notice of Allowance for U.S. Appl. No. 10/139,831, mailed Jun. 26, 2007, 25 pages.

Non-Final Office Action for U.S. Appl. No. 11/828,246, mailed Jun. 15, 2009, 26 pages.

Notice of Allowance for U.S. Appl. No. 11/828,246, mailed Nov. 16, 2009, 20 pages.

Non-Final Office Action for U.S. Appl. No. 10/140,088, mailed Apr. 27, 2006, 13 pages.

Notice of Allowance for U.S. Appl. No. 10/140,088, mailed Sep. 7, 2006, 13 pages.

Notice of Allowance for U.S. Appl. No. 10/140,088, mailed Oct. 24, 2006, 8 pages.

Notice of Allowance for U.S. Appl. No. 10/140,088, mailed Jan. 11, 2007, 5 pages.

Non-Final Office Action for U.S. Appl. No. 11/621,038, mailed Apr. 23, 2009, 44 pages.

Final Office Action for U.S. Appl. No. 11/621,038, mailed Dec. 23, 2009, 27 pages.

Non-Final Office Action for U.S. Appl. No. 12/198,697, mailed Feb. 2, 2010, 50 pages.

Non-Final Office Action for U.S. Appl. No. 10/140,749, mailed Aug. 10, 2006, 22 pages.

Final Office Action for U.S. Appl. No. 10/140,749, mailed Jun. 27, 2007, 23 pages.

Final Office Action for U.S. Appl. No. 10/140,749, mailed Jan. 8, 2008, 23 pages.

Non-Final Office Action for U.S. Appl. No. 10/140,749, mailed Jun. 6, 2008, 28 pages.

Final Office Action for U.S. Appl. No. 10/140,749, mailed Dec. 8, 2008, 30 pages.

Non-Final Office Action for U.S. Appl. No. 10/140,749, mailed May 27, 2009, 38 pages.

Final Office Action for U.S. Appl. No. 10/140,749, mailed Jan. 13, 2010, 44 pages.

Non-Final Office Action for U.S. Appl. No. 10/140,753, mailed Apr. 20, 2006, 11 pages.

Final Office Action for U.S. Appl. No. 10/140,753, mailed Jan. 10, 2007, 27 pages.

Non-Final Office Action for U.S. Appl. No. 10/140,753, mailed Aug. 22, 2007, 14 pages.

Non-Final Office Action for U.S. Appl. No. 10/140,753, mailed Jan. 8, 2008, 14 pages.

Final Office Action for U.S. Appl. No. 10/140,753, mailed Aug. 25, 2008, 22 pages.

Requirement for Restriction/Election for U.S. Appl. No. 11/000,359, mailed Jun. 20, 2008, 7 pages.

Non-Final Office Action for U.S. Appl. No. 11/000,359, mailed Oct. 23, 2008, 28 pages.

Non-Final Office Action for U.S. Appl. No. 11/000,359, mailed May 29, 2009, 14 pages.

Notice of Allowance for U.S. Appl. No. 11/000,359, mailed Sep. 22, 2009, 17 pages.

Requirement for Restriction/Election for U.S. Appl. No. 11/118,697, mailed Jun. 2, 2009, 8 pages.

Notice of Allowance for U.S. Appl. No. 11/118,697, mailed Sep. 30, 2009, 41 pages.

Non-Final Office Action for U.S. Appl. No. 09/855,038, mailed Jun. 2, 2005, 14 pages.

Final Office Action for U.S. Appl. No. 09/855,038, mailed Feb. 7, 2006, 8 pages.

Non-Final Office Action for U.S. Appl. No. 09/855,038, mailed Oct. 4, 2006, 14 pages.

Notice of Allowance for U.S. Appl. No. 09/855,038, mailed Apr. 26, 2007, 8 pages.

Requirement for Restriction/Election for U.S. Appl. No. 09/988,066, mailed Dec. 13, 2005, 7 pages.

Non-Final Office Action for U.S. Appl. No. 09/988,066, mailed Jul. 14, 2006, 17 pages.

Non-Final Office Action for U.S. Appl. No. 09/988,066, mailed Apr. 6, 2007, 22 pages.

Final Office Action for U.S. Appl. No. 09/988,066, mailed Oct. 31, 2007, 16 pages.

Advisory Action for U.S. Appl. No. 09/988,066, mailed May 28, 2008, 4 pages.

Notice of Allowance for U.S. Appl. No. 09/988,066, mailed Oct. 30, 2008, 16 pages.

Notice of Allowance for U.S. Appl. No. 09/988,066, Mailed Jan. 9, 2009.

Non Final Office Action U.S. Appl. No. 11/804,977, mailed Jan. 14, 2008, 13 pages.

Notice of Allowance for U.S. Appl. No. 11/804,977, mailed Nov. 19, 2008, 17 pages.

Non-Final Office Action for U.S. Appl. No. 09/855,015, mailed Oct. 28, 2004, 12 pages.

Non-Final Office Action for U.S. Appl. No. 09/855,015, mailed Jan. 12, 2006, 6 pages.

Notice of Allowance for U.S. Appl. No. 09/855,015, mailed Sep. 8, 2006, 3 pages.

Requirement for Restriction/Election for U.S. Appl. No. 09/855,015, mailed Nov. 3, 2006, 6 pages.

Notice of Allowance for U.S. Appl. No. 09/855,015, mailed Jan. 7, 2008, 4 pages.



(56)

**References Cited**

## OTHER PUBLICATIONS

Supplemental Notice of Allowance for U.S. Appl. No. 09/855,015, mailed Feb. 4, 2008, 3 pages.

Non-Final Office Action for U.S. Appl. No. 11/611,067, mailed Feb. 20, 2009, 11 pages.

Final Office Action for U.S. Appl. No. 11/611,067, mailed Oct. 16, 2009, 35 pages.

Non-Final Office Action for U.S. Appl. No. 11/611,067, mailed Dec. 8, 2009, 11 pages.

Non-Final Office Action for U.S. Appl. No. 11/615,769, mailed Apr. 15, 2009, 11 pages.

Final Office Action for U.S. Appl. No. 11/615,769, mailed Jan. 22, 2010, 34 pages.

U.S. Appl. No. Appl. 12/880,518, filed Sep. 13, 2010, Wong.

Non-Final Office Action for U.S. Appl. No. 10/140,751, mailed Dec. 20, 2010, 23 pages.

Notice of Allowance for U.S. Appl. No. 12/400,645, mailed on Jan. 26, 2011, 14 pages.

U.S. Appl. No. 13/862,160, filed on Apr. 12, 2013, by Ramanathan et al.

U.S. Appl. No. 13/925,564, filed on Jun. 24, 2013 by Wong et al.

U.S. Appl. No. 13/766,330, filed on Feb. 13, 2013 by Wong et al.

Non-Final Office Action for U.S. Appl. No. 12/608,985 mailed on May 31, 2013, 9 pages.

Notice of Allowance for U.S. Appl. No. 12/198,710 mailed on May 28, 2013, 10 pages.

Non-Final Office Action for U.S. Appl. No. 13/398,725 mailed on Aug. 30, 2013, 8 pages.

Notice of Allowance for U.S. Appl. No. 12/684,022 mailed on Aug. 20, 2013, 10 pages.

Notice of Allowance for U.S. Appl. No. 13/083,481 mailed on Sep. 3, 2013, 9 pages.

Non-Final Office Action for U.S. Appl. No. 10/832,086 mailed on Sep. 9, 2013, 13 pages.

Non-Final Office Action for U.S. Appl. No. 13/458,650 mailed on Oct. 2, 2013, 5 pages.

Notice of Allowance for U.S. Appl. No. 11/745,008 mailed on Oct. 7, 2013, 9 pages.

U.S. Appl. No. 12/795,492, filed Jun. 7, 2010, Davis et al.

U.S. Appl. No. 12/702,031, filed Feb. 8, 2010, Davis.

U.S. Appl. No. 12/466,277, filed May 14, 2009, Lin.

U.S. Appl. No. 12/198,710, filed Aug. 26, 2008, Zhang et al.

U.S. Appl. No. 12/883,073, filed Sep. 15, 2010, Davis.

U.S. Appl. No. 12/900,279, filed Oct. 7, 2010, Bansal et al.

Non-Final Office Action for U.S. Appl. No. 11/646,845, mailed on Oct. 4, 2010, 48 pages.

Final Office Action for U.S. Appl. No. 11/779,714, mailed on Nov. 9, 2010, 24 pages.

Non-Final Office Action for U.S. Appl. No. 12/198,697, mailed on Oct. 25, 2010, 36 pages.

Non-Final Office Action for U.S. Appl. No. 12/198,710, mailed on Sep. 28, 2010, 15 pages.

Requirement for Restriction/Election for U.S. Appl. No. 12/639,749, mailed on Dec. 7, 2010, 3 pages.

Non-Final Office Action for U.S. Appl. No. 12/639,762, mailed on Sep. 1, 2010, 40 pages.

Final Office Action for U.S. Appl. No. 12/400,594, mailed on Oct. 28, 2010, 13 pages.

Non-Final Office for U.S. Appl. No. 12/400,645, mailed on Sep. 1, 2010, 45 pages.

Non-Final Office Action for U.S. Appl. No. 12/372,390, mailed on Sep. 13, 2010, 10 pages.

Non-Final Office Action for U.S. Appl. No. 12/505,390, mailed on Oct. 28, 2010, 51 pages.

Final Office Action for U.S. Appl. No. 12/070,893, mailed on Nov. 24, 2010, 11 pages.

Non-Final Office Action for U.S. Appl. No. 13/862,160 mailed on Jun. 17, 2014, 11 pages.

\* cited by examiner






Fig. 1

10






Fig. 1a






Fig. 2






Fig. 3

318



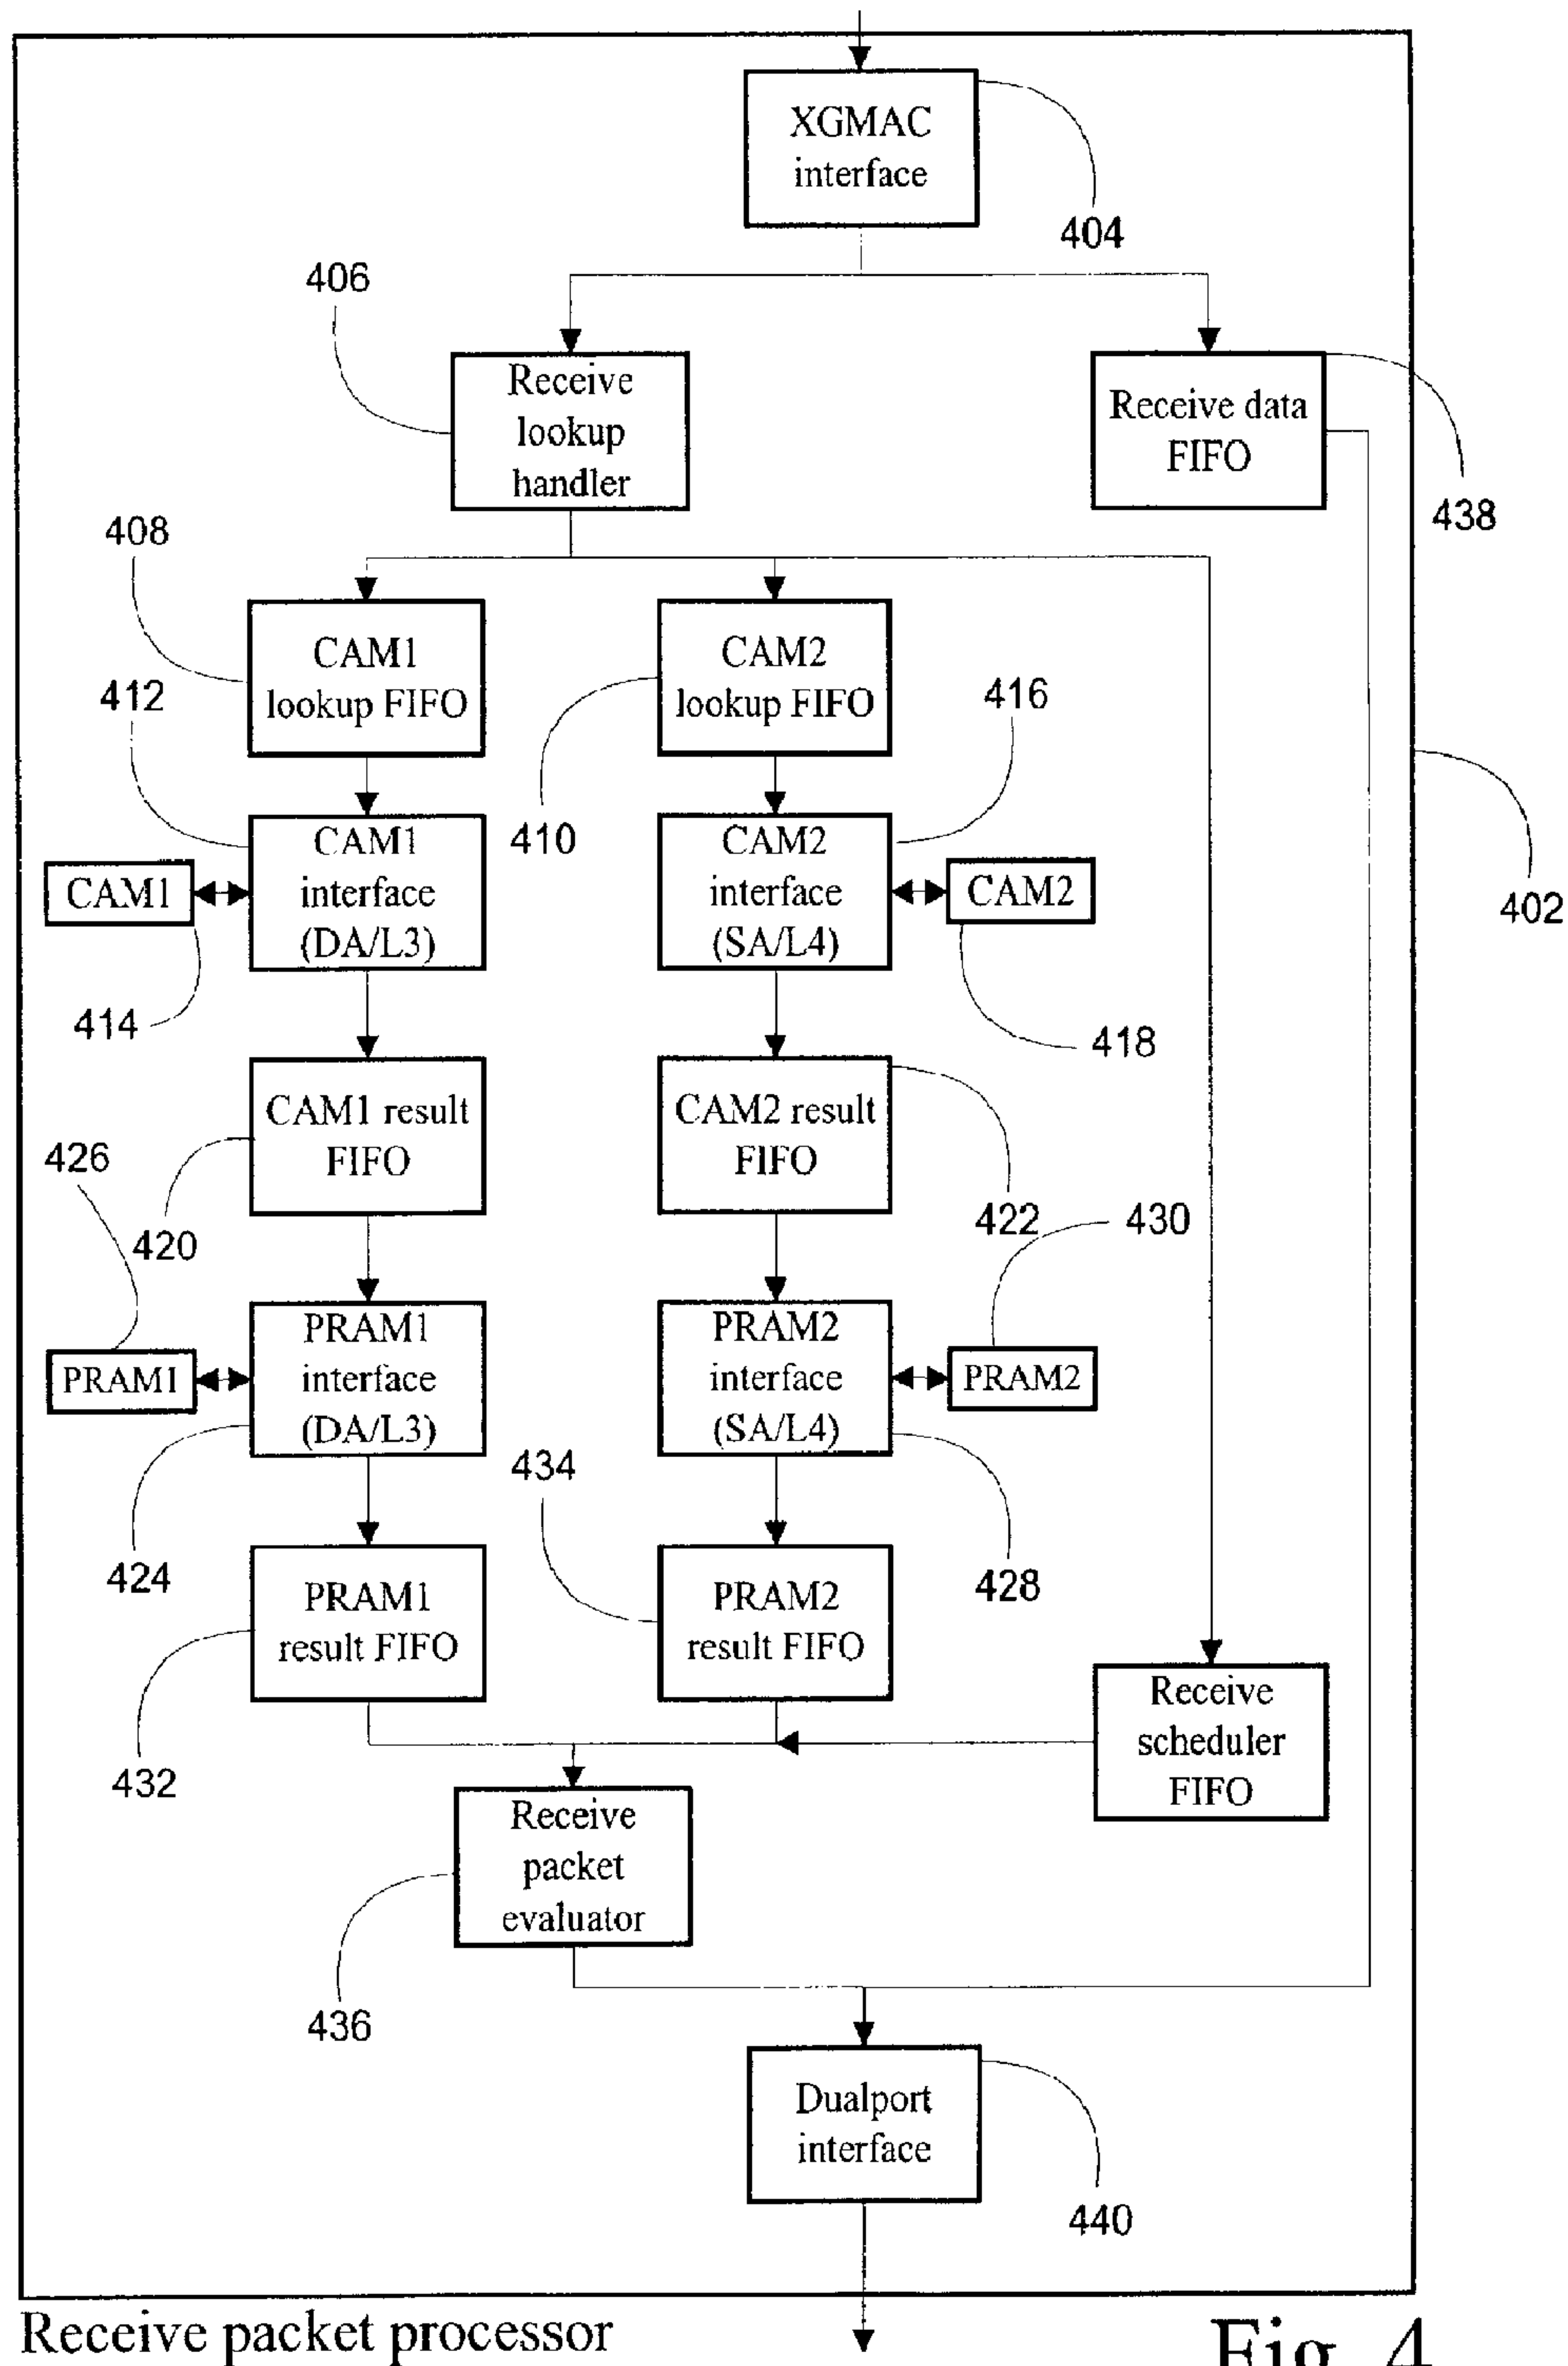


Fig. 4




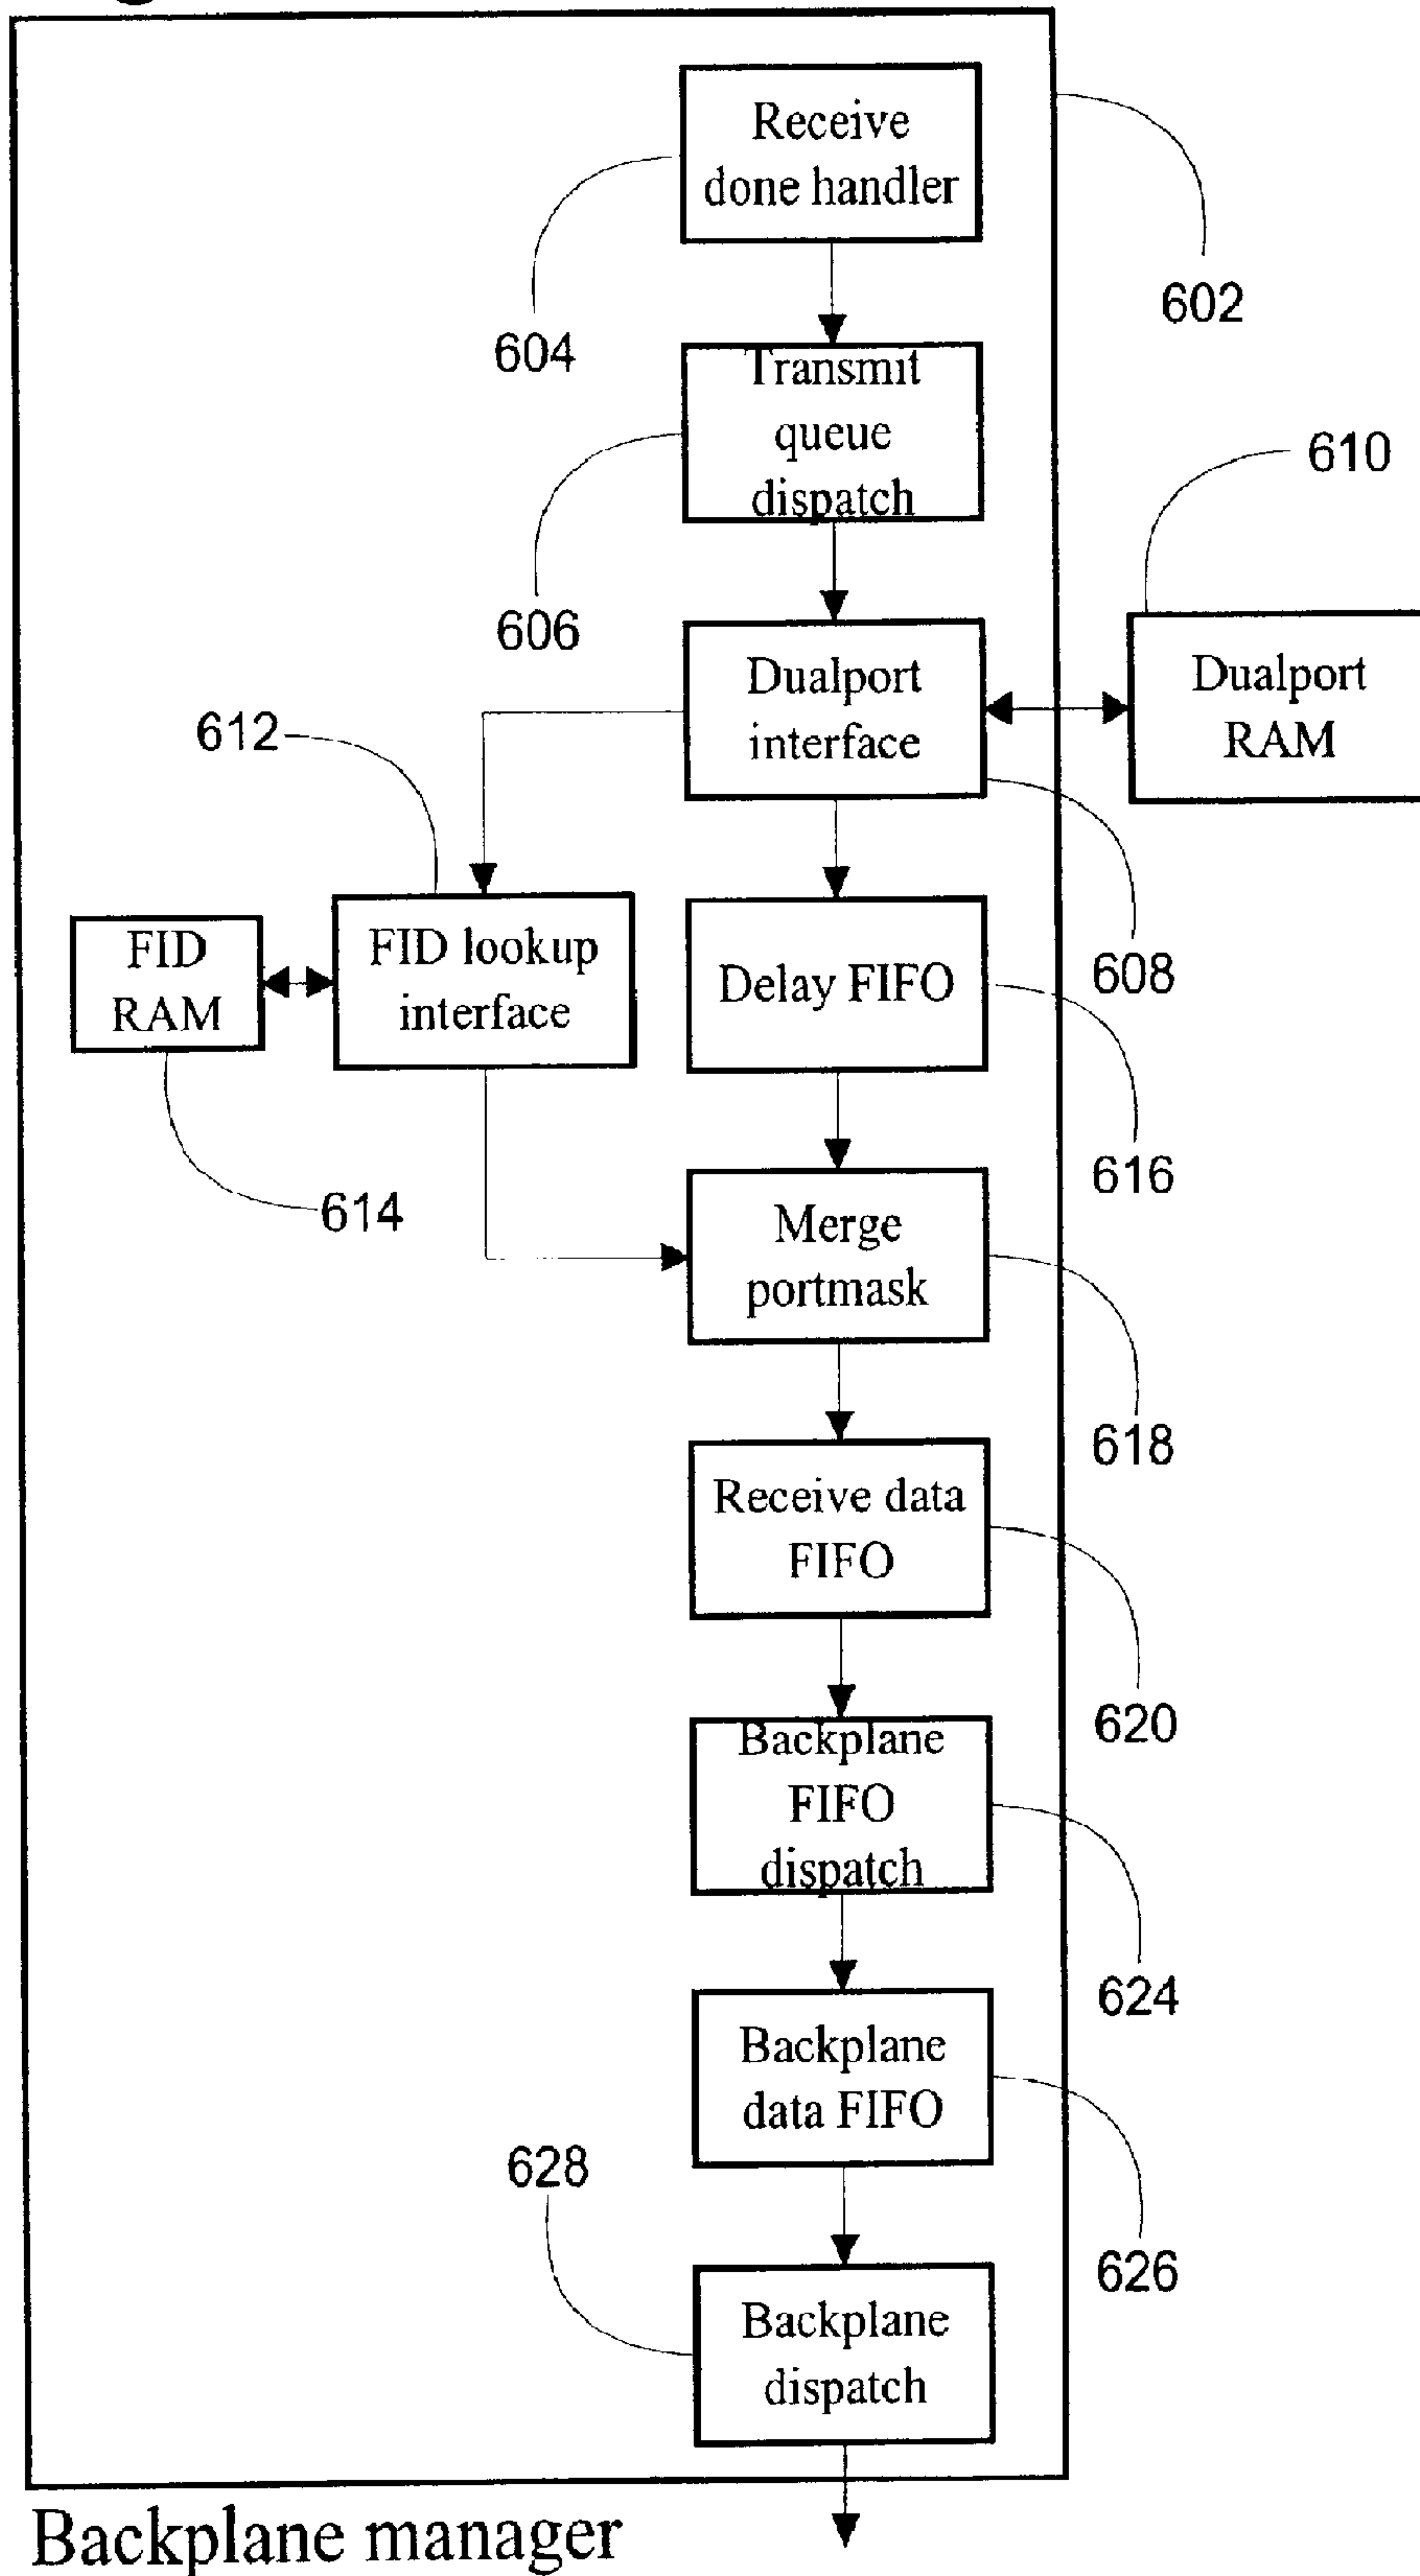


Fig. 5

Fig. 6








Fig. 7




Fig. 8





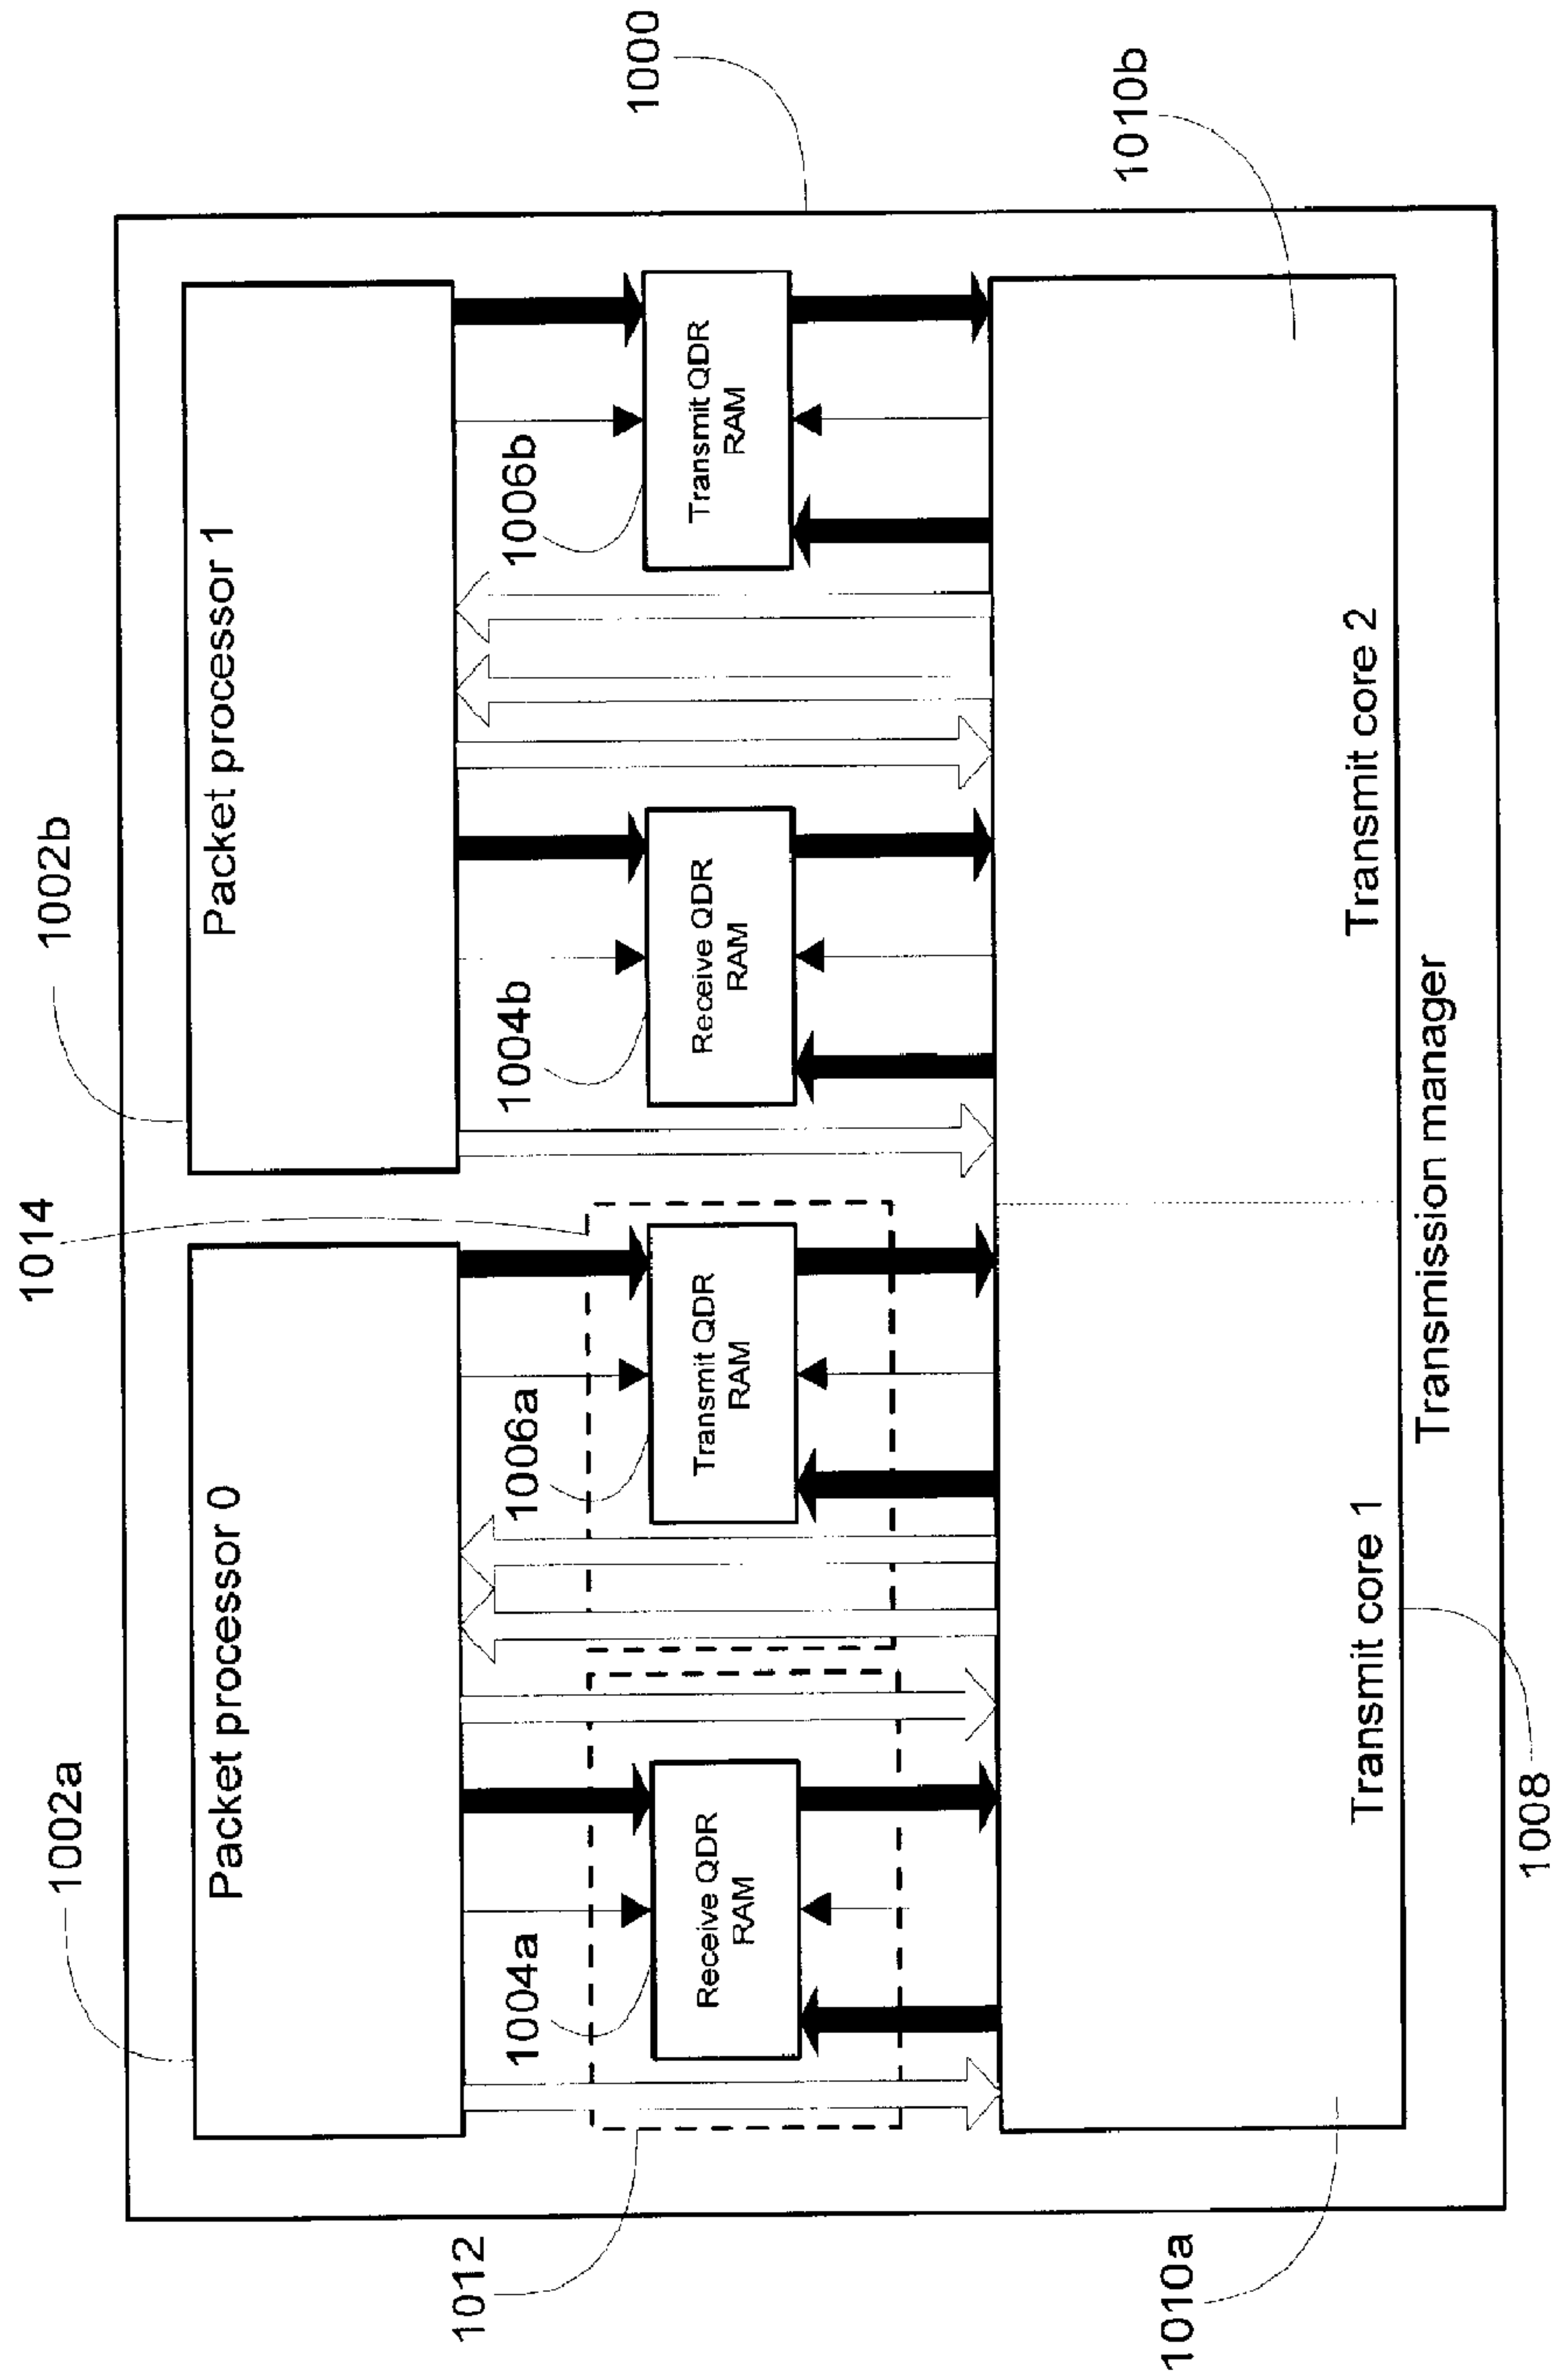


Fig. 10






Fig. 11




Fig. 12



## SYSTEM AND METHOD FOR HIGH SPEED PACKET TRANSMISSION

### CROSS-REFERENCE TO RELATED APPLICATIONS

This application is a divisional of U.S. application Ser. No. 11/000,359, filed Nov. 29, 2004, which is a divisional of U.S. application Ser. No. 10/438,545, filed May 15, 2003, now U.S. Pat. No. 6,901,072, all of which are incorporated herein by reference in their entirety for all purposes.

### COPYRIGHT NOTICE

A portion of the disclosure of this patent document contains material which is subject to copyright protection. The copyright owner has no objection to the facsimile reproduction by anyone of the patent document or the patent disclosure, as it appears in the Patent and Trademark Office patent files or records, but otherwise reserves all copyright rights whatsoever.

### BACKGROUND OF THE INVENTION

The invention described herein relates to computer networking and, in particular, to improved methods, systems, and software for routing data at high speeds through a switch or other network routing device.

The explosive growth of the Internet has brought more and more users online every day, and computer networks have assumed an increasingly important role in today's highly interconnected world. As users increasingly rely on the network to deliver required data, network traffic has increased exponentially. Moreover, with the adoption of new and more bandwidth-intensive applications, enormous burdens are placed on network infrastructure. Network administrators are thus constantly seeking faster and more reliable methods and equipment to transport data to accommodate these demands.

Ethernet, one of the earliest networking protocols, is today the most widely used method to transport data in computer networks. Robert Metcalf and David Boggs developed Ethernet as an experiment at the XEROX Palo Alto Research Center in 1973. At Ethernet's inception, the struggle to accommodate users' needs for bandwidth had not yet started. As network traffic demands at this time were quite low, Ethernet initially had a data transmission rate of 2.94 megabits per second (Mbps).

Metcalf, however, recognized the potential for rapid network growth and posited a theorem now known as "Metcalf's Law" which states that the value of a network expands exponentially as the number of users increase. Gordon Moore, an expert in the field of semi-conductor development, posited another theorem known as Moore's Law which states that the power of microprocessors will double every 18 months and their price will be reduced by half. When taken together, these two laws predict rapid growth of networking technologies: as users join the network, more people will want to join at an exponential rate equivalent to the rise in value of the network, while processing technologies to support this growth and faster transport are constantly increasing at rapidly diminishing costs.

The evolution of Ethernet technologies has followed theory. The first commercial release of Ethernet occurred in 1979 with a transmission rate of 10 Mbps—more than a three-fold increase over the experimental system created just five years earlier. Ethernet went through a variety of standardizations during the 1980s and line rates remained constant at

10 Mbps while the technology matured. In 1995, however, Ethernet became available at 100 Mbps. In 1998, bandwidth jumped again to 1 gigabit per second (Gbps). Most recently, a new standard was adopted for Ethernet transmission rates at 10 Gbps representing a 100-fold increase in seven years.

Implementation of 10 Gbps network infrastructure requires overcoming significant hurdles not addressed by current advances in the art. For example, previous generations of Ethernet technology, although fast, had an ample number of clocks in which to perform packet analysis and retransmit data. With the rise of 10 Gbps Ethernet, however, calculations previously carried out over a given number of clocks must now be completed in a fraction of the time so that the desired bandwidth is in fact available.

There is thus a need for a systems and methods capable of efficiently accommodating data transfer rates over a network in excess of 10 Gbps.

### SUMMARY OF THE INVENTION

The present invention provides systems and methods for providing data transmission speeds at or in excess of 10 gigabits per second between one or more source devices and one or more destination devices. According to one embodiment, the system of the present invention comprises a first and second media access control (MAC) interfaces to facilitate receipt and transmission of packets over an associated set of physical interfaces. The system also contemplates a first and second field programmable gate arrays (FPGA) coupled to the MAC interfaces and an associated first and second memory structures, the first and second FPGAs are configured to perform initial processing of packets received from the first and second MAC interfaces and to schedule the transmission of packets to the first and second MAC interface for transmission to one or more destination devices. The first and second FPGAs are further operative to dispatch and retrieve packets to and from the first and second memory structures. A third FPGA, coupled to the first and second memory structures and a backplane, is operative to retrieve and dispatch packets to and from the first and second memory structures, compute appropriate destinations for packets and organize packets for transmission. The third FPGA is further operative to receive and dispatch packets to and from the backplane.

In accordance with certain aspects of the present invention, the third FPGA is an integrated circuit which is configured to retrieve and dispatch packets to and from one or more memory structures in a switching or routing device and a backplane in the switching or routing device. The integrated circuit includes a plurality of cores each operative to provide a receive and transmit pipeline for packets corresponding to a given one of the ports on the switching or routing device. The first and second transmit cores are further operative to receive packets from the memory structures and process the packets for dispatch to their intended destinations. The IC also includes a local switching FIFO in each core operative to transfer packets between the core and a second of the cores corresponding to a port for the second core.

In some embodiments, the cores each comprise a transmit quality of service module operative to determine the order in which packets received at a given one of the cores from the backplane and the local switching FIFO are transmitted to a given one of the memory structures. This order may be determined based on quality of service data read from the header of the packet and in accordance with any desired prioritization algorithm.



In further embodiments, the integrated circuit has a backplane transmit sorter, coupled between the cores and the backplane, operative to organize packets for dispatch to one or more slots on the backplane. The backplane transmit sorter may implement a number of FIFOs to organize packets prior to transmission to the one or more slots on the backplane, which may be equal to a number of slots on the backplane.

The integrated circuit may further include other elements, such as a receive FIFO arbiter, coupled between the cores and the backplane, operative to arbitrate the receipt of data from a plurality of transmit cores for placement in one or more receive FIFOs, and a backplane transmit grouper, coupled between the cores and the backplane, operative to modify the size of a packet for transmission over the backplane.

In accordance with other aspects, the present invention, a system is described for providing data transmission speeds at or in excess of 10 gigabits per second between one or more source devices and one or more destination devices. The system includes first and second transmit and receive data pipelines, the first and second transmit and receive data pipelines being used to couple first and second media access control (MAC) interfaces to facilitate receipt and transmission packets over a physical interface, and first and second packet processors configured to perform initial processing of received packets that are buffered in first and second memory structures. The system also includes third and fourth transmit and receive data pipelines used to couple a transmission manager to a backplane and the first and second transmit and receive data pipelines. The third and fourth transmit and receive data pipelines are configured to receive packets from a backplane, organize the packets for transmission that are buffered in the first and second memory structures and schedule the transmission of packets on the backplane.

In accordance with still further aspects of the present invention, a method is described for providing data transmission speeds at or in excess of 10 gigabits per second between one or more source devices and one or more destination devices. The method involves receiving packets over a first and second media access control (MAC) interfaces, the first and second MAC interfaces operative to facilitate receipt and transmission of packets over a first and second physical interfaces. Data is pipelined bi-directionally through first and second processors coupled to the first and second MAC interfaces and first and second memory structures. The first and second processors are each configured to perform initial processing of received packets to be buffered in the first and second memory structures and scheduling packets for transmission to the MAC interfaces for transmission to one or more destination devices over the first and second physical interfaces. The first and second processors are further operative to dispatch and retrieve packets to and from the first and second memory structures. Dual bi-directional data pipelines are implemented through a third processor coupled to the first and second memory structures and a backplane, the second processor configured to compute an appropriate destination for a packet and organize packets for transmission.

#### BRIEF DESCRIPTION OF THE DRAWINGS

The invention is illustrated in the figures of the accompanying drawings which are meant to be exemplary and not limiting, in which like references are intended to refer to like or corresponding parts, and in which:

FIG. 1 is a block diagram of a system architecture for an Ethernet blade in accordance with one embodiment of the present invention;

FIG. 1A is a block diagram of a system architecture for an Ethernet blade in accordance with a second embodiment of the present invention;

FIG. 2 is a high level flow diagram of a connection of a packet processor component of the present invention to an outside network, in accordance with one embodiment of the present invention;

FIG. 3 is a block diagram of receive and transmit packet processors of one embodiment of the present invention;

FIG. 4 is a block diagram of a receive packet processor in accordance with one embodiment of the present invention;

FIG. 5 is a flow diagram showing the data flow in the receive packet processor of FIG. 4 in accordance with one embodiment of the present invention;

FIG. 6 is a block diagram of a backplane manager in accordance with one embodiment of the present invention;

FIG. 7 is a flow diagram showing the data flow in a transmission accumulator in accordance with one embodiment of the present invention; and

FIG. 8 is a block diagram of a transmit packet processor component in accordance with one embodiment of the present invention;

FIG. 9 is a block diagram of a system architecture for an Ethernet blade in accordance with a third embodiment of the present invention;

FIG. 10 is a block diagram of an exemplary data flow between two packet processors and a transmission manager in an Ethernet blade in accordance with one embodiment of the present invention;

FIG. 11 is a block diagram presenting the components comprising a transmission manager and its associated external structures in accordance with one embodiment of the present invention; and

FIG. 12 is a block diagram presenting a FID port mask used assist in to properly routing packets within an Ethernet blade in accordance with one embodiment of the present invention.

#### DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

Embodiments of methods and systems according to the present invention are described through reference to FIGS. 1 through 8. Turning to FIG. 1, a block diagram is presented depicting a high-level schematic of the components of one possible embodiment of the invention to allow data transfer speeds at or in excess of 10 gigabits per second. As shown, the invention comprises a printed circuit board ("PCB") 10 used to house and provide interconnections for a media access controller ("MAC") 12, a packet processor ("PP") 14, one or more content addressable memory ("CAM") controllers 16, one or more controllers for random access memories containing parameter information ("PRAM") processors 18, a receive dual-port memory buffer 20, a transmit dual-port memory buffer 22, a transmission manager 24, and a backplane interface 26.

The PCB 10 provides a surface on which to place other components of the invention. The PCB 10, also known as a "blade" or "module", can be inserted into a slot on the chassis of a network traffic management device such as a switch or a router. This modular design allows for flexible configurations with different combinations of blades in the various slots of the device according to differing network topologies and switching requirements. Furthermore, additional ports for increased network connectivity may easily added by plugging additional blades into free slots located in the chassis.

An example of such a switch is the BigIron® switch produced by Foundry Networks, Inc. of San Jose, Calif. The



BigIron switch chassis consists of multiple distributed switching modules each of which contain a high-bandwidth memory system for scalable chassis bandwidth. The local switching fabric of the BigIron switch houses the forwarding engines, provides packet-level examination and classification based on Layer 2/3/4 information, and performs IP subnet look-ups and packet modifications of IP and IPX packets.

The MAC 12 is the interface by which data is received and transmitted to and from the network. In one embodiment, such network data comprises Ethernet packets. The MAC 12 forwards received packets to the PP 14 for further processing and also receives packets for transmission to the network from the PP 14. The MAC 12 performs any data conversions required for network data to be processed by the PP 14 for routing within the device chassis and for data processed by PP 14 to be transmitted to the network. For example, in one embodiment of the invention, the MAC 12 performs data conversions because network data comprises 32 bit double data rate (“DDR”) data, while the PP 14 processes only 64 bit single data rate (“SRD”) data. The MAC is typically responsible for data validity checking, as well as data gathering.

The PP 14 is a processor chip responsible for receiving packets from the MAC 12 and processing them for forwarding through the device chassis, as well as for processing packets received from the device chassis intended for transmission over the network. These two functions, while performed on the same chip, are preferably performed simultaneously and in parallel. There are thus, in a sense, two pipelines in the PP 14: a receive pipeline for processing network packets intended for transmission within the chassis and a transmit pipeline for processing internally routed packets intended for transmission over the network.

In one embodiment of the invention, the packet processor is a field programmable gate array (“FPGA”), which is an integrated circuit that can be programmed in the field after manufacture. An advantage of using FPGAs with the invention is that an FPGA provides significant flexibility over an application specific integrated circuit (“ASIC”) and is also much less expensive to prototype and implement.

The receive pipeline of the PP 14 is responsible for packet classification, performing CAM and PRAM lookups, generating packet headers for forwarding packets through a chassis, and preparing packet modifications. Network packets are received by the PP 14 from the MAC 12 in multi-byte bursts based on scheduling priorities determined at the MAC 12. The PP 14 examines packets and extracts packet forwarding information from the packets such as the destination address (“DA”) of the packet and the source address (“SA”) of the packet. The PP 14 extracts the type of service (“TOS”), whether the packet has a virtual local area network (“VLAN”) tag, session related data such as in the case of IPv4 or IPX data, and other additional Layer 3 and Layer 4 information useful in routing the packet through the chassis. The PP 14 passes this forwarding information extracted from the packet header to a CAM processor 16 for further processing.

The CAM controller or processor 16 takes information forwarded by the PP 14 and performs a lookup comparing this information to data stored in a local memory of the CAM processor 16. If the information matches information stored in the local memory of the CAM processor 16, additional forwarding information regarding disposition of the packet is available in the local memory of the PRAM processor 18 and can be retrieved for future incorporation into the packet header.

When such successful CAM matches occur, the PRAM processor 18 retrieves additional forwarding information from its local memory for incorporation into the header of the

packet. The packet is reformatted with a new internal hardware header for routing the packet within the chassis and stored in the receive dual-port memory buffer 20 for processing by the transmission manager. This internal hardware header is also sometimes referred to as a chassis header.

An important technique in implementing the invention is pipelining. Pipelining is an advanced technique used by processors, wherein a processor begins executing a subsequent instruction before a prior instruction has finished executing. Accordingly, a processor can have multiple instructions processing in its “pipeline” simultaneously with each instruction at a different processing stage.

The pipeline is divided into processing segments, with each segment executing its operation concurrently with the other segments. When a segment completes its operation, it passes the result to the next segment in the pipeline and fetches data for processing from the preceding segment. Often, temporary memory buffers are used to hold data values between segments, which allows operations to complete faster since each segment no longer waits for the other segment to finish processing prior to handing off data. The final results of the process emerge at the end of the pipeline in rapid succession.

The receive dual-port memory 20 (as well as its counterpart, the transmit dual-port memory 22) acts as a pipeline buffer in the embodiment of the invention depicted in FIG. 1. The receive dual-port memory 20 enables the PP 14 to store processed data and continue processing the next packet without having to wait for the transmission manager 24 to become available, thereby expediting operations of both the PP 14 and the transmission manager 24. Other buffers are used throughout the invention and in its various components to achieve pipelining and faster packet processing in an analogous manner.

The transmit pipeline of the PP 14 retrieves data from the transmit dual-port memory 22 according to a programmable priority scheme. The PP 14 extracts network destinations from the dual-port data and reassembles packet header forwarding information by removing any packet header modifications that take place in order to route the packet through the switch chassis. The PP 14 performs sanity checks on packet data to ensure that only those packets intended for transmission are passed on to the MAC 12.

Since packets routed through the chassis carry header information pertaining to forwarding within the chassis, this information must be removed and replaced with header forwarding information appropriate for routing over the network. After the proper network header forwarding information is reassembled and the chassis header information is removed, the PP 14 forwards the data to the MAC 12 for eventual transmission over the network to the intended address.

While the PP 14 handles traffic to and from the MAC 12 and conversions of packet headers between network packet headers and internal chassis packet headers, the transmission manager 24 handles traffic flow to and from the backplane interface 114. Like the PP 14, the transmission manager 24 is a processor chip that implements a dual pipeline architecture: a receive pipeline for network data to be internally routed within the device chassis and a transmit pipeline for internally routed data intended for network transmission. These two functions, while performed on the same chip, are preferably performed in parallel according to one embodiment of the invention. In one embodiment of the invention, the transmission manager 24 is an FPGA, although use of other processor types is within the scope of the invention.



The transmission manager **24** fetches network data intended for routing through the device chassis from the receive dual-port memory **20** and stores internally routed data intended for network transmission in the transmit dual-port memory **22**. The receive pipeline of the transmission manager **24** retrieves data from the receive dual-port memory **20** according to instructions issued to the transmission manager **24** by the PP **14**. The transmission manager **24** determines data transmission priority for the data retrieved and schedules transmissions to the backplane **26** according to this priority scheme. In one embodiment of the invention, there are four different priority levels assigned to data.

The transmission manager **24** extracts backplane destinations from data, and sends data to those destinations according to predetermined priority algorithms. Backplane destinations may comprise other blades in the chassis or, in some cases, may comprise the blade of the transmission manager **24** itself, which is called “one-armed routing.”

The transmit pipeline of the transmission manager **24** handles internally routed packets received from the backplane interface **26** and intended for transmission over the network. The transmission manager **24** collects packets from the backplane interface **26** and organizes them into per-source, per-priority transmit queues stored in the transmit dual-port memory **22**. The transmission manager **24** notifies the PP **14** when a packet is stored in the transmit dual-port memory **22** and available for processing.

FIG. **1a**, presents a block diagram depicting a high-level schematic of the components of an alternative embodiment of the invention. As shown, the invention comprises a printed circuit board **100**, a media access controller **102**, a receive packet processor **104** (“RXPP”), one or more CAM processors **106**, one or more PRAM memory processors **108**, a receive dual-port memory buffer **110**, a backplane manager **112**, a backplane interface **114**, a transmission accumulator (“TX accumulator”) **116**, a transmit dual-port memory buffer **118**, and a transmit packet processor (“TXPP”) **120**.

The PCB **100** provides a surface on which to place many of the other components of the invention. The PCB **100**, also known as a “blade” or “module”, can be inserted into one of a plurality of slots on the chassis of a network traffic management device such as a switch or a router. This modular design allows for flexible configurations with different combinations of blades in the various slots of the device according to differing network topologies and switching requirements.

The MAC **102** is the interface by which a blade receives and transmits data to and from the network. In one embodiment, such network data comprises Ethernet packets. The MAC **102** forwards received packets to the RXPP **104** for further processing and receives packets for transmission to the network from the TXPP **120**. The MAC **102** also performs any data conversions required for network data to be processed by the RXPP **104** or for data processed by TXPP **120** to be transmitted to the network. For example, the MAC **102** may perform data timing conversions where network data comprises 32 bit DDR data while the RXPP **104** and the TXPP **120** process only 64 bit SDR data. The receive packet processor **104** is responsible for packet classification, performing CAM and PRAM lookups, generating packet headers for forwarding packets through a chassis, and preparing packet modifications. In one embodiment of the invention, the receive packet processor **104** is an FPGA. In an alternate embodiment of the invention, the RXPP **104** is an ASIC. Packets are received by the RXPP **104** from the MAC **102** in multi-byte bursts based on scheduling priorities determined at the MAC **102**. The RXPP **104** examines packets and extracts packet forwarding information from a packet, such as

the destination address of the packet and the source address of the packet. The RXPP **104** extracts the TOS, any defined VLAN tags, session related data such as in the case of Ipv4 or TX data, and other additional Layer 3 and Layer 4 information useful in routing the packet through the chassis. The RXPP **104** passes this forwarding information to one of the CAM processors **106** for further examination.

The CAM processor **106** takes information forwarded by the RXPP **104** and performs a lookup, comparing received information to data stored in local memory of the CAM processor **106**. If the comparison returns a match, additional forwarding information regarding disposition of the packet is stored in local memory of one of the PRAM processors **108** and can be retrieved for future incorporation into the packet header. The PRAM processor **108** retrieves additional forwarding information from its local memory for incorporation into the header of packet. The packet is then stored in the receive dual-port memory buffer **110** for processing by the backplane manager **112**. Those of skill in the art will recognize that additional processing may be performed before storage in the receive dual port memory.

The receive dual-port memory **110** (as well as its counterpart, the transmit dual-port memory **118**) acts as a pipeline buffer between processes. The receive dual-port memory **110** enables the RXPP **104** to store processed data and continue processing the next packet without having to wait for the backplane manager **112** to become available. Pipelining operation execution expedites processing of both the RXPP **104** and the backplane manager **112**. Other buffers are used throughout the invention and within its various components to achieve pipelining and faster packet processing in this manner.

The next segment in the receive pipeline is the backplane manager **112**. The backplane manager **112** is a processor designed for retrieving data from the receive dual-port memory buffer **110** and dispatching packets to the backplane interface **114**. Data is retrieved from the receive dual-port memory **110** according to instructions issued to the backplane manager **112** by the RXPP **104**. The backplane manager **112** determines data transmission priority for the data retrieved and schedules transmissions to the backplane **114** according to this priority scheme. According to one embodiment of the invention, there are four different priority levels assigned to data.

The backplane manager **112** extracts backplane destinations from received data; the data sent to indicated destinations according to programmable priority algorithms. Backplane destinations may comprise other blades in the chassis or, in the case of OAR, may comprise the blade of the backplane manager **112** that initially receives the data. When packets scheduled for OAR are detected, they are forwarded to the transmission accumulator **116** via the OAR data path as shown in FIG. **1a**. In one embodiment of the invention, the backplane manager **112** is an FPGA. In an alternate embodiment of the invention, the backplane manager **112** is an ASIC.

The transmit accumulator **116** is a processor that receives packet data from the backplane **114** intended for transmission. The transmit accumulator **116** collects packets from the backplane **114** and organizes them into per-backplane-source, per-priority transmit queues stored in the transmit dual-port memory **118**. The transmit accumulator **116** notifies the TXPP **120** when data comprising a packet is stored in the transmit dual-port memory **118** and available for processing. In one embodiment of the invention, the transmit accumulator **116** is an FPGA.

The transmit packet processor **120** retrieves data from the transmit dual-port memory **118** according to a programmable



priority scheme. The TXPP 120 extracts network destinations from the data and reassembles packet header forwarding information by removing any packet header modifications that took place in order to route the packet through the device chassis. The TXPP 120 performs sanity checks on packet data to ensure that only those packets intended for transmission are passed on to the MAC 102. Since packets routed through the chassis carry header information pertaining to forwarding within the chassis, this information must be removed and replaced with header forwarding information appropriate for routing over the network. After the proper network header forwarding information is reassembled and the chassis header information is removed, the transmit packet processor 120 forwards the data to the MAC 102 for eventual transmission over the network to the intended address. In one embodiment of the invention, the transmit packet processor 120 is an FPGA. In an alternate embodiment of the invention, the transmit packet processor 120 is an ASIC.

FIG. 2 presents a high-level schematic of one embodiment of the invention as it connects to a network, e.g., an optical network comprising fiber optic connections. The optics block 202 is the interface through which all network data traffic passes. The optics block 202 contains a transmitter for generating the optical signals to the network when data is received from the transceiver 204. In some embodiments, the transmitter might comprise a laser or a light emitting diode. The optics block 202 also contains a detector for receiving optical data traffic from the network. When optical data is received, a photodetector generates an electrical current that is amplified to level useable by the transceiver 204. The signal is then communicated to the transceiver 204 for further processing.

The transceiver 204 directs the transmission and receipt of signals to and from the optics block 202. The transceiver 204 receives electrical data signals intended for transmission to the MAC 206 and instructs the transmitter in the optics block 202 to generate optical signals corresponding to the electrical data signals. Conversely, the transceiver 204 receives electrical data signals from the optics block 202 and passes these signals to the MAC 206 for processing.

There are many asynchronous boundaries between the various components of the invention. For example, data passes to and from the transceiver 204 and the MAC 206 at a fixed speed. In one embodiment of the invention, the datapath 208 between the transceiver and the MAC 206 operates sending 4 clock signals along with 32 bit DDR data at 156.25 MHz. The datapath 212 between the MAC 206 and the packet processor 210, however, may operate at a different speed. For example, in one embodiment of the present invention, the datapath 212 between the MAC 206 and the packet processor 210 operates sending 4 clock signals along with 64-bit SDR at 66 MHz. Multiple clock signals are sent with the data and used to minimize timing differences between groups of data signals and a clock. In one embodiment of the invention, one clock signal is included per 8 bits of DDR data and one clock signal is included per 16 bits of SDR data. In addition to clock signals, control signals are also sent along with data to indicate packet boundaries and possible error conditions. In one embodiment of the invention, control signals are distributed across 4 clock groups of data.

Those skilled in the art will recognize that an important technique in managing the dataflow between these asynchronous boundaries is the use of FIFO buffers that permit the dataflow to remain synchronized. Given the extremely high rate of data transfer provided by the invention, conventional techniques for clock distribution, such as those known in the art and used in the case of personal computer boards, will not

allow reliable capture and transfer of data between components of the invention operating according to different clocks. The invention, therefore, implements source synchronous clocking wherein the clock is sent along with the data.

When the clock arrives at the packet processor 210 from the MAC 206, for example, the clock is exactly in relationship according to the MAC 206, but the packet processor 210 can also capture the data on that clock via a FIFO. Data from the MAC 206 is captured inside a FIFO, which allows the packet processor to synchronize, in the presence of this data, between the source synchronous clock contained in the FIFO data and the clock the packet processor 210 is using at its core.

The invention uses source synchronous clocking in a symmetric manner. For example, data passing from the packet processor 210 to the MAC 206 is also captured in a FIFO to allow the MAC 206 to synchronize, in the presence of the FIFO data, between the source synchronous clock (of the packet processor 210 core) and the clock the MAC 206 is using at its core clock. In an alternative embodiment, the invention also implements differential source synchronous clocking which is known to those skilled in the art. Differential source synchronous clocking works in much the same manner as source synchronous clocking, except that two clock signals are sent with the data instead of one clock signal. The two clock signals, a high and low signal, are used to calculate a more precise approximation of the signal value being transmitted which those skilled in the art will recognize is used to reduce noise and generate more accurate data transmissions.

FIG. 3 is a block diagram depicting one embodiment of the components of the MAC 102 as presented in FIGS. 1 and 1a. Components of the MAC 102 are embodied in the MAC processor chip 302. According to one embodiment of the invention, the MAC chip 302 is an FPGA. In an alternate embodiment of the invention, the MAC chip 302 is an ASIC. The MAC 102 is the interface between the network via the PHY transceiver 300 and the RXPP 104 and TXPP 120 packet processor chips. According to one embodiment of the invention, the MAC 102 communicates directly to the PHY layer transceiver 300 via a DDR interface and with the packet processor chips of the RXPP 104 and the TXPP 120 via an SDR interface.

The PHY transceiver 300 is the component applying signals to the network wire and detecting signals passing through the network wire. According to one preferred embodiment of the invention, the PHY transceiver 300 is a 10 Gigabit Ethernet transceiver transmitting and receiving 32 bit DDR data at 156.25 Mhz. Data received by the PHY transceiver 300 is passed to the receive front end 306 of the MAC 102. The receive front end 306 is an interface that receives data, which is passed to the receive block 304 for further processing. According to one preferred embodiment of the invention, the receive front end 306 receives 32 bit DDR data.

The receive block 304 performs a variety of tasks on data received from the receive front end 306 and is very flexible in operation. The receive block 304 internally converts data received from the receive front end 306 into a format suitable for transmission to the RXPP 104. According to one embodiment of the invention, the receive block converts 32 bit DDR data into 64 bit SDR data for transmission. The receive block 304 may also perform other tasks as required according to various embodiments of the invention such as verifying and extracting XGMII tokens, realigning bytes such that the start of packet ("SOP") token is placed in a "lane zero" position, verifying SOP and EOP framing, detecting giant packets,



verifying and optionally stripping packet cyclic redundancy checks, tracking the full suite of RMON statistics, and other useful operations.

The receive block **304** also generates flow control packets via the pause and flow control sync block **332**. The receive block **304** operates off of the recovered source synchronous clocks contained in the incoming data packets received from the PHY transceiver **300**. Other components of the MAC **102**, including the transmit block **328**, however, are operating off of an internal core clock generated locally. Although these two clocks are nominally the same frequency, there is some variance since they are not really the same clock and therefore tend to “drift” over time. This difference between the two clocks requires periodic synchronization of the receive block **304** and the transmit block **328** for the purposes of passing flow control messages to generate pause frames and avoid network congestion.

In such a scenario, the receive block **304** receives an incoming message from a remote source (to which the transmit block **328** is sending data) indicating that the remote source is becoming congested and requesting that the transmit block **328** pause transmission for a requested interval. The pause and flow control sync block **332** synchronizes the receive block **304** clock with the transmit block **328** clock to permit the receive block **304** to pass the pause frame request to the transmit block **328** and reduce the network congestion. Conversely, in the unlikely event that the receive FIFO RAM **308** becomes congested, the pause and flow control sync block **332** would synchronize the two clocks to permit the receive block **304** to instruct the transmit block **328** to start issuing flow control pause frames to a remote sender to reduce network congestion in the MAC **102**.

The receive block **304** passes processed data to the receive FIFO RAM **308** via the write port **310** of the receive FIFO RAM **308** which enables the receive block **304** to process the next packet without waiting for the receive FIFO block **314** to become available. The receive FIFO RAM **308** is a two-port memory having a write port **310** that accepts incoming data from the receive block **304** and a read port **312** that transmits data stored in the receive FIFO RAM **308** to the receive FIFO block **314**. The write port **310** and the read port **312** operate independently of each other thus permitting more efficient use of the receive FIFO RAM **308** by the receive block **304** and the receive FIFO block **314**.

The FIFO RAM **308** further permits data flow through the asynchronous boundary. In one embodiment of the invention, the receive block **304** operates at a different speed than the receive FIFO block **314**. Thus, the FIFO RAM **308** acts as a bridge, allowing data flow to be synchronized between these asynchronous components. For example, in the Foundry BigIron switch, the receive block **304** operates at a 156.25 MHz clock recovered from the arriving data and the FIFO block **314** operates on a locally generated 156.25 MHz clock that differs slightly and drifts in phase relationship over time.

To further reduce processing time, the receive block **304** starts streaming data into the receive FIFO RAM **308** when the receive block detects the start of a packet and stops streaming data into the receive FIFO RAM **308** when the receive block **304** detects the end of the packet. All of the packet processing components of the invention stream data into FIFOs in this manner which greatly reduces processing time since components are not required to wait until an entire packet is finished processing to start copying the packet into a FIFO.

The receive FIFO block **314** reads data stored in the receive FIFO RAM **308** via the read port **312**. The receive FIFO block **314** also notifies the RXPP **104** that packet data is contained

in the receive FIFO RAM **308** and available for transmission. This data is transmitted to the RXPP **104** for further processing. According to one embodiment of the invention, the receive block FIFO **314** transmits 64 bit SDR data to the RXPP **104**.

In addition to the receive pipeline of the MAC **102** as set forth above, the MAC **102** also contains a transmit pipeline that operates in a similar fashion with similar flexibility. The transmit FIFO block **320** is the interface of the MAC **102** that receives data from the TXPP **120**. According to one embodiment of the invention, the transmit FIFO block **320** receives 64 bit SDR data from the TXPP **120**.

The transmit FIFO block **320** streams received data to the transmit FIFO RAM **322** via the write port **324** of the transmit FIFO RAM **322**, enabling the transmit FIFO block **320** to process the next incoming packet without waiting for the transmit block **328** to become available. The transmit FIFO RAM **322** is a two-port memory having a write port **324** that accepts incoming data from the transmit FIFO block **320** and a read port **326** that transmits data stored in the transmit FIFO RAM **322** to the transmit block **328**. Similar to the two-port memory comprising the receive FIFO RAM **308**, the write port **324** and the read port **326** of the transmit FIFO RAM **322** operate independently of each other, thus permitting pipelining and more efficient use of the transmit FIFO RAM **322** by the transmit FIFO block **320** and the transmit block **328**.

The transmit block **328** reads data stored in the transmit FIFO RAM **322** via the read port **326**. Similar to the receive block **304**, the transmit block **328** performs a variety of tasks and is very flexible in operation. The transmit block **328** internally converts data received from TXPP **120** into a format suitable for transmission to the PHY transceiver **300**. According to one embodiment of the invention, the transmit block converts 64 bit SDR data into 32 bit DDR data for transmission. The transmit FIFO RAM **322** facilitates this conversion by bridging the asynchronous boundary between the transmit block **328** and the transmit FIFO block **320**.

The transmit block performs other tasks as required according to embodiments of the invention, such as generating flow control packets to the PHY side sender at the request of the TXPP **120** (and in addition to internal flow control requests generated by the receive block **304** via the pause and flow control sync **332** when the receive FIFO RAM **308** is full) to avoid network congestion, calculating and optionally appending a cyclic redundancy check to a packet, determining and inserting XGMII tokens, and tracking the full suite of RMON statistics. In one embodiment of the invention, the transmit block **328** stores data in a programmable FIFO buffer used for data rate matching which allows the MAC **102** to connect to a packet processor that is receiving data slower than line rate.

The transmit block **328** passes data processed for to the transmit front end **330** thus enabling the transmit block **328** to begin processing the next packet. The transmit front end **330** is an interface that receives data from the transmit block **328** and passes this data to the PHY transceiver **300** for transmission over the network. According to one preferred embodiment of the invention, the transmit front end **330** transmits 32 bit DDR data to the PHY transceiver **300**.

Building on the illustration presented in FIG. 1, FIG. 4 presents a block diagram depicting one embodiment of the components of the RXPP **104**. The RXPP **402** is responsible for packet classification, performing CAM and PRAM look-ups, generating hardware packet headers used for internally forwarding packets within the chassis of the network device, such as between blades, and for preparing packet modifications. Components of the RXPP **104** are embodied in the



RXPP chip **402**. According to one preferred embodiment of the invention, the RXPP chip **402** comprises an FPGA. In an alternate embodiment of the invention, the RXPP chip **402** comprises an ASIC.

The XGMAC **404** interface is responsible for requesting data for the RXPP **402** from the MAC **102**. When the receive lookup handler **406** is available to parse additional data and the receive data FIFO **438** is available to store additional data, the XGMAC **404** instructs the MAC **102** to begin streaming packet data into the RXPP **104**. The XGMAC interface **404** is connected to the MAC **102** and receives data for processing by the RXPP **104**. The XGMAC interface **404** also acts as an asynchronous boundary, writing source-synchronous 64-bit data from the MAC **102** in a small internal FIFO, then sending the synchronized data at 66 MHz in 256-bit chunks for subsequent processing.

The XGMAC interface **404** sends synchronized data as it is received from the MAC **102** to the receive data FIFO **438**, where it is held until CAM and PRAM lookups are performed. The receive data FIFO **438** thus acts as a delay buffer until packet processing is completed and the packet data can start being written by the dual-port interface **440** into the receive dual-port memory **110**.

While all data related to a packet is streamed to the receive data FIFO **438**, the XGMAC interface **404** also parses the incoming data as it is received from the MAC **102** and streams only the packet header information to the receive lookup handler **406** where it will be used to perform CAM and PRAM lookups.

The receive lookup handler **406** performs sanity checks on the packet data as it is received from the XGMAC interface **404**. For example, the receive lookup handler **406** identifies valid packet contexts by identifying consistent start-of-packet and end-of-packet boundaries. In this respect, the receive lookup handler **406** also monitors a bad packet control signal from the MAC **102** indicating a data fault. If a data fault is detected, the receive lookup handler **406** discards the header data from the bad packet and also flushes any associated data already stored in the receive data FIFO **438** related to the bad packet. In one embodiment of the invention, if packet processing has already started, a data fault flag indicating a bad packet is stored in the receive data FIFO **438**. The dual port interface **440** will later discard the packet when the data fault flag is retrieved from the receive data FIFO **438**.

The receive lookup handler **406** strips VLAN tags, compares the packet MAC destination address against the port MAC address, performs IPv4 TOS field lookups as required, and also checks the protocol used to encode the packet. Examples of encoding protocols include IP, IP ARP, IPv4, IPv6, 802.3, IPX RAW, IPX LLC, IPX 8137, IPX SNAP, Appletalk, Appletalk ARP, NetBios, IP SNAP, and IP ARP SNAP. This information will be used to assemble an internal hardware packet header to be appended to the packet for use in forwarding the data internally throughout the chassis of the network switch. This additional information is passed from the receive lookup handler **406** to the RX scheduler FIFO **407**. The RX scheduler FIFO **407** holds this information until the CAM and PRAM lookups are completed on the destination and source addresses extracted by the receive lookup handler **406** from the packet header.

Based upon the information extracted, the receive lookup handler **406** forms the CAM lookups and builds part of the hardware packet header for internally forwarding the packet through the chassis of the network device. The internal state of the receive lookup handler **406** containing this information is then split into two CAM lookup FIFOs **408** and **410**, which are memory buffers that permit the receive lookup handler

**406** to start processing the next packet received from the XGMAC interface **404**. Packet processing is thus pipelined, allowing the receive lookup processor **406** to continue processing packets without waiting for either the CAM1 interface **412** or the CAM2 interface **410** to become available. Information relating to the destination address of the packet and other protocol fields from the header related to Layer 3 are passed to CAM1 lookup FIFO **408**. Information relating to the source address of the packet and other protocol fields from the header related to Layer 4 are passed to CAM2 lookup FIFO **410**. In an alternate embodiment of the invention, the two pipelines are merged into a single pipeline containing a single CAM interface and a single FIFO interface for lookups.

The CAM1 interface **412** becomes available, retrieves the data stored in the CAM1 lookup FIFO **408**, and submits requests regarding this data to the external ternary CAM1 **414** memory bank that contains a data array of values against which to perform lookups. The CAM1 interface **412** is also pipelined and supports dispatching lookups for multiple packets to the external ternary CAM1 **414** memory bank since it takes longer than four clocks for the external CAM1 **414** to respond.

If the lookup generates a match against an entry in the CAM1 **414** array, additional forwarding information exists in the PRAM1 **426** memory bank regarding the disposition of the packet. Forwarding information might include details such as the destination port of the packet, the port mirror requirement, the packet type, VLAN handling information, packet prioritization data, multicast group membership, replacement destination MAC addresses (used in network routing), and/or other similar packet data known in the art. The CAM1 **414** array entry also contains a link to the memory address of the additional forwarding information stored in the PRAM1 **426** memory bank. This link is stored in the CAM1 result FIFO **420** until the PRAM1 interface **424** is available to perform lookups.

Similarly, the CAM2 interface **416** retrieves source address data from the CAM2 lookup FIFO **410**, performs lookups by submitting requests to the external ternary CAM2 memory bank **418**, and stores the results of these lookups in the CAM2 result FIFO **422** until the PRAM2 interface **428** is available to perform lookups. According to one embodiment of the invention, the CAM2 interface **416** operates in parallel with the CAM1 interface **412** to allow CAM lookup operations to complete faster.

The PRAM1 interface **424** retrieves the data associated with the successful CAM1 interface **412** lookups from the CAM1 result FIFO **420**. The PRAM1 interface **424** extracts from this data the link to the memory address of the additional forwarding information stored in the PRAM1 **426** memory bank. PRAM1 interface **424** lookup results are stored in the PRAM1 result FIFO so work can immediately start on the next packet. According to one embodiment, PRAM lookups for a packet take 3 clocks. Similarly, and preferably in parallel, the PRAM2 interface **428** retrieves data associated with successful CAM2 interface **416** source address lookups from the CAM2 result FIFO **422**, performs lookups to obtain additional forwarding information stored in the PRAM2 **430** memory bank, and stores the results in the PRAM2 result FIFO **434**.

The receive packet evaluator **436** extracts the data from the PRAM2 result FIFO **432**, PRAM2 result **434**, and the RX scheduler FIFO **407**. The receive packet evaluator **436** uses this information to construct the internal hardware header used to forward a packet through the chassis with the most advanced forwarding in this aspect permitting total destina-



tion address/VLAN/TOS replacement and packet header modification to support hardware packet routing. In one embodiment of the invention, the internal hardware header comprises sixteen bytes. The receive packet evaluator **436** also determines the priority level of the packet according to the CAM and PRAM lookups and may optionally adjust the packet priority according to whether the packet is VLAN tagged or contains IPv4 TOS fields. The priority level is inserted into the internal hardware header of the packet.

The receive packet evaluator **436** notifies the dual-port interface **440** that processing is complete and passes the new internal hardware header to the dual-port interface **440** for integration with the packet data stored in the receive data FIFO **438**. The dual-port interface **440** reads from the receive data FIFO **438**, applying packet modifications to incorporate the new hardware packet header and stores this packet data in the receive dual-port memory **110**. The dual-port interface **440** also detects the end of packet (“EOP”) signal and issues a receive packet processing completion notification to the backplane manager **112** so the backplane manager **112** will know to retrieve the packet. If a packet is flagged as bad (for example, an invalid cyclic redundancy check) the buffer is instead immediately recycled for the next packet and the current packet is deleted.

FIG. 5 presents a block diagram depicting the operations of the RXPP **402** presented in FIG. 4 more discretely. Data flow commences with the receive lookup handler **501** receiving packet data from the XGMAC interface **404** as illustrated in FIG. 4. The XGMAC interface **404** parses data received from the MAC **102** and sends only the packet header information to the receive lookup handler **501**.

The receive port tracker **502** examines the port information contained in the packet header to ensure that any VLAN information tags contained in the packet header will be accepted at the destination address port. If the destination address port is not configured to accept the packet header VLAN information or lack thereof, then the receive lookup handler **501** either sets an error bit in the packet header if debugging is supported or the packet is discarded.

Alternatively, the receive lookup handler **501** will strip the VLAN tag from its field in the packet and store the VLAN tag in the internal hardware packet header for future use.

The receive lookup handler **501** checks the protocol used to encode the packet and classifies the packet accordingly in block **504**. Examples of encoding protocols include IP, IP ARP, IPv4, IPv6, 802.3, IPX RAW, IPX LLC, IPX 8137, IPX SNAP, Appletalk, Appletalk ARP, NetBios, IP SNAP, and IP ARP SNAP. This information is used to assemble an internal hardware packet header to be appended to the packet for use in forwarding the data internally throughout the chassis of the switch. This additional information is passed from the receive lookup handler **501** to the RX scheduler FIFO **522**. The RX scheduler FIFO **522** holds this information until the CAM and PRAM lookups are completed on the destination and source addresses extracted by the receive lookup handler **501** from the packet header.

The receive lookup handler **501** also forms the CAM lookups and builds part of the hardware packet header in block **506**. The receive lookup handler **501** extracts source and destination address information from the packet header for use in the CAM lookups. The internal state of the receive lookup processor **501** containing this information is then passed to the CAM lookup FIFO **508**, which is a memory buffer that permits the receive lookup processor **501** to start processing the next packet received from the XGMAC interface **404**. Packet processing is thus pipelined allowing the

receive lookup processor **501** to continue efficiently processing packets without waiting for the CAM interface **509** to become available.

When the CAM interface **509** becomes available, it fetches the address data stored in the CAM lookup FIFO **508** as shown in block **510**. The CAM interface **509** dispatches requests regarding data in block **512** to the external ternary CAM memory **516** that contains a data array of values against which to perform lookups. The CAM interface **509** is pipelined and supports cycling lookups for multiple packets to the external ternary CAM **516** memory since it takes longer than four clocks for the external CAM **516** to respond. Block **514** illustrates a programmable delay incorporated into the CAM interface **509** pipeline that compensates for this delay while the CAM lookup is being performed.

If the lookup generates a match against an entry in the CAM array **516**, additional forwarding information regarding disposition of the packet is available in the PRAM memory **530**. Forwarding information might include details such as the destination port of the packet, the port mirror requirement, the packet type, VLAN handling information, packet prioritization data, multicast group membership, and/or other similar packet data known in the art. The CAM array **516** entry also contains a link to the memory address of the additional forwarding information stored in the PRAM memory **530**. This link is returned by the CAM memory **516** as shown in block **518** and stored in the PRAM lookup FIFO **520** until the PRAM interface **523** is available to perform lookups.

When the PRAM interface **523** becomes available, it fetches the link to the address in the PRAM memory **530** that is stored in the PRAM lookup FIFO **520** as shown in block **524**. In block **526**, the PRAM interface **523** dispatches requests to retrieve the additional forwarding information for the packet to the external PRAM memory **530**. The PRAM interface **523** is pipelined and supports cycling lookups for multiple packets to the external PRAM memory **530** since it takes multiple clocks for the external PRAM memory **530** to return results from a lookup. Block **528** illustrates a programmable delay incorporated into the PRAM interface **523** pipeline that compensates for this delay while the PRAM lookup is being performed. The external PRAM **530** returns the additional forwarding information in block **532** and these results are stored in the PRAM result FIFO **534** until the receive packet evaluator **535** is available.

In block **536**, the receive packet evaluator **535** fetches data from the PRAM result FIFO **534** and the receive scheduler FIFO **522**. The receive packet evaluator **535** evaluates this information in block **538** and uses the results to construct the internal hardware packet header in block **540**. The internal hardware packet header is used to forward the packet through the chassis among other blades inserted into slots on the backplane. The most advanced forwarding in this aspect permits total destination address/VLAN/TOS replacement and packet header modification to support hardware packet routing. In one embodiment of the invention, the internal hardware header comprises sixteen bytes.

The receive packet evaluator **535** notifies the dual-port interface **542** that processing is complete and passes the new internal hardware header to the dual-port interface **542** for integration with the packet data stored in the receive data FIFO **438**, as illustrated in FIG. 4. The dual-port interface **542** reads from the receive data FIFO **438** applying packet modifications to incorporate the new hardware packet header for internally forwarding the packet through the chassis of the switch and stores this packet data in the receive dual-port memory **110**. The receive dual-port memory is organized as four large FIFOs corresponding to four exemplary priority



levels. The dual-port interface **440** also detects the end of packet (“EOP”) and issues a receive packet processing completion notification to the backplane manager **112** so the backplane manager **112** will know to retrieve the packet. If a packet is flagged as bad (for example, an invalid cyclic redundancy check) the packet is deleted and the buffer is recycled for the next packet.

Transport within a blade continues with FIG. **6**, which presents a block diagram depicting the components of the backplane manager **112** as illustrated in FIG. **1**. Components of the backplane manager **602** are embodied in the backplane manager chip. According to an embodiment of the invention, the backplane manager chip **602** comprises an FPGA.

The backplane manager **602** is responsible for retrieving data from the receive dual-port memory **610**, determining backplane destinations for this data, and sending this data to those destinations. The backplane manager **112** also manages four large FIFOs stored in the external dual-port memory **610**. These FIFOs store data according to priority levels by which the data is to be processed by the backplane manager **112**.

The receive done handler **604** receives EOP information from the receive packet processor **104**, including information regarding packet length and packet priority. This information is used to assist the receive done handler **604** in tracking receive dual-port memory **110** utilization for the four priority levels and scheduling packets for dispatch by the transmit queue dispatch **606**. If the backplane manager **602** or the receive dual-port memory FIFOs **610** are running low on resources, the receive done handler **604** sends a throttle control back to the receive packet processor **104**.

The transmit queue dispatch **606** is responsible for ordered packet dispatch from the four priority levels of the receive dual-port memory FIFOs **610**. The transmit queue dispatch **606** receives packet length and priority information from the receive done handler **604** and uses this information to schedule packet retrieval from the dual-port RAM **610** by the dual-port interface **608** according to prioritization algorithms contained in the transmit queue dispatch **606**.

According to one embodiment of the invention, absolute priority is used with higher priority packets being unconditionally transmitted before any packets of lower priority. Absolute priority, however, is not always desirable. In another embodiment, some fraction of the transmission bandwidth available to the backplane manager **112** is dedicated to lower priority packet transmission regardless of whether higher priority packets are also pending because packets are often received by the invention faster than they can be transmitted. If some bandwidth were not allocated to lower priority packets in this manner, a bottleneck might be created with lower priority packets not being transmitted due to higher priority packets monopolizing all available transmission bandwidth. Packets are thus scheduled and posted for use by the transmit queue dispatch **606**.

The dual-port interface **608** fetches data from the receive dual-port memory **610** based on instructions received by the transmit queue dispatch **606**. At the start-of-packet boundary, the dual-port interface **608** extracts a forwarding identifier (“FID”) from the packet and sends the FID to the FID lookup interface **612**. The FID is an abstract chassis/system wide number used to forward packets. Each packet type has a FID to instruct the blade how to handle a given type of packet. This allows each blade in the chassis to look at the FID separately to decide how to individually forward the packet.

The FID lookup interface **612** translates the FID received from the dual-port interface **608** into a port mask by performing a lookup against addresses stored in the external FID RAM **614**. The port mask is a multi-bit field representing a

port on the blade and also other possible backplane slot destinations in the device chassis. According to one embodiment, the port mask is an 8-bit field representing a 10 Gigabit Ethernet port on the blade and seven other possible backplane slot destinations.

The FID lookup takes a number of clock cycles to complete during which time read data is posted to the delay FIFO **616** by the dual-port interface **608**. According to one embodiment of the invention, the FID lookup by the FID lookup interface **612** into the external FID RAM **614** requires a delay of six clocks to complete in order to resume processing the data.

The FID lookup is complete and the results are passed from the FID lookup interface **612** to the merge port mask **618**. Read data stored in the delay FIFO **616** is also passed to the merge port mask **618**. The merge port mask **618** integrates the read data with the appropriate FID lookup port mask result and other port masks as set forth below to ensure that the data is transmitted to all intended destinations.

The merge port mask **618** takes the FID lookup port mask result and combines it with CPU and monitor information stored in configuration registers of the backplane manager. For example, a FID indicates a physical destination or possibly a list of destinations, but the receive packet processor **104** might have determined that the CPU also needs a copy of the data and therefore sets the CPU flag for combination with the FID lookup port mask by the merge port mask **618**. Alternatively, when a packet needs to be sent to a monitor port for network debugging or similar purpose, the monitor port mask is combined with the FID port mask. The merge port mask **618** thus generates a “qualified” port mask indicating all destinations for which the packet data is intended.

The merge port mask **618** may also apply source port suppression. In certain situations, the blade that receives the data packet is listed as part of a FID port mask; source port suppression conditionally prevents the blade from retransmitting packets it just received. For example, this might occur in a broadcast situation where packets with unknown addresses are sent to all ports. Once all port mask data is combined with packet data, the merge port mask **618** stores the final result in the receive data FIFO **620** enabling the merge port mask **618** to process the next packet without waiting for the backplane FIFO dispatch **624** to become available.

The backplane FIFO dispatch **624** reads data from the receive data FIFO **620**, duplicating the data for each destination indicated in the qualified port mask. The backplane FIFO dispatch **624** restructures the data into a format required by the backplane, generates backplane state and slot information, and posts the results into the backplane data FIFO **626**. The backplane data FIFO **626** also acts as an asynchronous boundary between the backplane manager **602** core clock and the actual backplane clock. By posting the results in the backplane data FIFO **626**, the backplane FIFO dispatch **624** can process the next packet without waiting for the backplane dispatch **628** to become available. In one embodiment of the invention, data posted to the backplane data FIFO **626** is equivalent to two backplane transfers since the backplane manager runs at approximately one-half the clock speed of the backplane interface **114**.

The backplane dispatch **628** reads data from the backplane data FIFO **626** and outputs the data to the backplane via the backplane interface **114**. According to one embodiment, the backplane dispatch **628** reads data from the backplane data FIFO **626** suitable for more than one transfer because the ratio of the backplane interface **114** clock speed and the clock speed of the backplane manager **602** is not identical. In such an embodiment, the backplane dispatch **628** reads the number of transfers from the backplane data FIFO **626** that fully



utilizes the transmission capacity of the backplane interface **114**. For example, if the clock speed of the backplane interface **114** is double that of the backplane manager **602**, then the backplane dispatch **628** will read two transfers from the backplane data FIFO.

The backplane dispatch **628** also monitors backplane status and directs backplane transmission rates since it is possible for a backplane slot destination to become congested or otherwise unavailable. For example, if a plurality of blades comprising a single chassis are devoting all of their transmission capacities to a single blade, then they may overload the destination blade. Such a case might occur when two blades both transmit at 8 Gbps to a single destination blade that, according to the capacity of a backplane slot, can only receive 8 Gbps it total. The two blades would have to throttle back transmissions to the destination blade to 4 Gbps to avoid congestion.

Data is received from the backplane by the transmission accumulator **116** as presented in FIG. 1. Turning to FIG. 7, the transmission accumulator **116** collects packets from the backplane and organizes them into per-source, per priority transmit FIFOs stored in the transmit dual-port memory **118**. Components of the transmission accumulator are embodied in the transmission accumulator chip **702**. According to one embodiment of the invention, the transmission accumulator chip **702** comprises an FPGA.

Data is received from the backplane by the backplane front end **704**. The backplane front end passes received data to the backplane slot receive accumulator **706**. The backplane slot receive accumulator **706** is divided into a series of equal storage structures or memory buffers, with one buffer allocated for each slot or source on the chassis of the device. According to one embodiment of the invention, the backplane slot receive accumulator **706** is divided into eight buffers for receipt of data.

When a particular quantity of data is received into one of the backplane slot receive accumulator **706** buffers, the backplane slot receive accumulator **706** notifies the backplane data polling logic **708** to indicate the buffer and priority of the data being stored. In one embodiment of the invention, the backplane slot receive accumulator **706** waits to notify the backplane data polling logic **708** until 32 bytes of data have been received in a bucket and transfers between the two components thus comprise 32 bytes. If the backplane slot receive accumulator **706** is full, then the transmission accumulator is congested and no longer accepts data until the congestion is relieved.

The backplane data polling logic **708** reads data from the backplane slot receive accumulator **706** and organizes data according to source and priority. If packets are aborted from the backplane, the backplane data polling logic **708** deletes the packet in order to avoid propagation of the packet to the TXPP **120**.

The backplane data polling logic **708** processes the data and the final result is stored in the backplane receive FIFO **710**, enabling the backplane data polling logic **708** to process the next packet without waiting for the dual-port interface **712** to become available. The backplane receive FIFO **710** also permits dataflow through the asynchronous boundary between the backplane data polling logic block **708** and the dual-port interface **712**.

The dual-port interface **712** reads data from the backplane receive FIFO **710** and stores this packet data in the transmit dual-port memory **118**. The dual-port interface **712** also detects valid end-of-packet (“EOP”) indications and notifies the TXPP **120** via transmission of an EOP message that a packet is available in the transmit dual-port memory **118**. The transmit dual-port memory **118** also comprises a series of

FIFOs similar to the receive dual-port memory **110**. Instead of only four total FIFOs, however, the transmit dual-port memory **118** has four FIFOs for each buffer of the backplane slot accumulator **706**, thereby comprising 28 FIFOs for these buffers, plus an additional four FIFOs for the OAR path, yielding a total of 32 FIFOs.

Transmission continues in FIG. 8, which depicts a block diagram of the components of the transmit packet processor **120** as illustrated in FIG. 1a. Components of the TXPP **120** are embodied in the TXPP chip **800**. According to an embodiment of the invention, the TXPP chip **800** comprises an FPGA. The TXPP **800** is responsible for retrieving data from the transmit dual-port memory **803**, determining network destinations for this data and sending data to identified destinations. The TXPP **120** strips hardware header forwarding information used to route packets throughout the chassis of the switch and replaces this information with header forwarding information necessary to route packets over the network. The TXPP **120** also manages the FIFOs priority queues stored in the transmit dual-port memory **803**. These FIFOs store data according to priority levels by which the data is to be processed by the TXPP **800**.

The transmit done handler **801** receives EOP information from the TX accumulator **116**, including information regarding packet length and packet priority. This information is used to assist the transmit done handler **801** in tracking transmit dual-port memory **803** utilization for the four priority levels and scheduling packets for dispatch in the transmit queue dispatch **802**. The transmit done handler **801** notifies the transmit queue dispatch **802** regarding packet availability and priority.

The transmit queue dispatch **802** is responsible for ordered packet retrieval and dispatch from the four priority levels of the transmit dual-port memory **803** FIFOs. According to one embodiment of the invention, absolute priority is used with higher priority packets being unconditionally transmitted before any packets of lower priority. Absolute priority, however, is not always desirable. In alternative embodiments, some fraction of the transmission bandwidth available to the TXPP **120** is dedicated to lower priority packet transmission regardless of whether higher priority packets are also pending because packets are often received by the invention faster than they can be transmitted. If some bandwidth were not allocated to lower priority packets in this manner, a bottleneck might be created with lower priority packets not being transmitted due to higher priority packets monopolizing all available transmission bandwidth. Packets are thus scheduled and posted for use by the dual-port handler **804**.

The dual-port handler **804** fetches the data from the transmit dual-port memory **803** according to instructions received from the transmit queue dispatch **802**. At the start-of-packet boundary, the dual-port handler **804** extracts the FID from the packet and sends the FID to the FID lookup block **808**. The dual-port handler **804** also extracts any VLAN tags from the packet and sends this information to the multicast start offset lookup block **806**.

In the FID lookup block **808**, the FID received from the dual-port handler **804** is used to perform a lookup against a FID table. The FID lookup block **808** functions similarly to the interaction between the FID lookup interface **612** and the FID RAM **614** as presented in FIG. 6. Accordingly, the results obtained from the FID table indicate how the packet should be handled for transmission by the receiving blade. For example, the FID might indicate that although the packet may have arrived at the blade, the packet should not be transmitted by the blade. This might occur in a broadcast situation where a packet is broadcast to all blades within a chassis. If the FID



lookup block **808** determines that a packet has been erroneously received in this manner, the packet is deleted and no longer processed by the TXPP **120**. In this sense, the FID lookup block **808** also functions as a transmit filter to ensure that only valid packets are actually sent out over the network.

Results of the FID lookup are stored in the delay FIFO **810**. This permits the FID lookup block **808** to begin processing the next packet without waiting for the context track and internal header removal block **814** to become available. Pipelining processing data in this manner allows packet processing operations by the TXPP **120** to complete faster.

While the FID lookup block **808** is processing the FID data the multicast start offset lookup block **806** is processing any VLAN tags received from the dual-port handler **804**. A VLAN is a local area network identifier that maps locations based on a basis other than physical location. For example, devices attached to a VLAN might be grouped according to department, division, application, etc. Devices that are part of the same VLAN behave as if they were connected to the same wire even though they may actually be physically connected to different segments of a LAN. VLANs are configured using software protocols rather than in hardware and are therefore extremely flexible with respect to implementation. For example, a computer may be moved to a different physical location on the same VLAN without any hardware reconfiguration.

VLAN tags placed in a header field indicate whether a packet is intended for routing over a VLAN. Additionally, the VLAN tag in the header may also indicate that a packet is intended for VLAN multicasting. VLAN multicasting occurs when a packet is sent over a VLAN to more than one destination address. Since the header of each packet must be changed to reflect each destination address during VLAN multicasting, this process can be very resource intensive when performed using software.

The multicast start offset lookup block **806** supports hardware VLAN multicast replication. The multicast start offset lookup block **806** examines the VLAN tag extracted from the packet header and performs a lookup against a table stored in RAM in the multicast start offset lookup block **806**. If the packet VLAN tag matches an entry in the table, additional information pertaining to that VLAN is available at an address location in a memory array stored in the multicast replacement lookup block **812**. For example, multicast replacement lookup block **812** might contain information to assist with setting unique VLAN ID values, VLAN priorities, and TXA/SAS/srcport suppression behaviors for each packet transmitted over the VLAN.

The multicast start offset lookup block **806** takes the address to the memory array location of the multicast replacement lookup block **812** and stores this result in the delay FIFO **810**. This permits the multicast start offset lookup block **806** to begin processing the next packet without waiting for the context track and internal header removal block **814** to become available. Pipelining processing in this manner allows packet processing operations by the TXPP **120** to complete faster.

In addition to enabling pipelining, the delay FIFO **810** also stores values from the FID lookup block **808** and the multicast start offset lookup block **806** for retrieval by the multicast replacement lookup block **812** and the context track and internal header removal block **814**. The multicast replacement lookup block **812** retrieves the results of the multicast start offset lookup block **806** calculations from the delay FIFO **810** for processing packets subject to VLAN routing. The multicast replacement lookup block **812** takes the address of the memory array location contained in the multi-

cast replacement lookup block **812** and retrieves the additional information that is stored at that location pertaining to routing over the VLAN tag referenced in the packet header. This information is passed to the context track and internal header removal block **814** for incorporation into the outgoing packet header.

Taking the results from the delay FIFO **810** and the multicast replacement lookup block **812**, the context track and internal header removal block **814** removes the internal hardware header from the packet and begins the process of assembling an outgoing packet header suitable for transmission over the network. Those skilled in the art will recognize that a number of manipulations to the outgoing packet header must take place before this can occur. The context track and internal header removal block **814** passes information regarding any data offset to the header which may have occurred to the barrel shifter **816**. The context track and internal header removal block **814** passes information regarding the TXA/PTYPE to the SA substitution and L3 assist block **818**. The context track and internal header removal block **814** passes information regarding the packet VLAN ID and the VLAN tag status to the VLAN insertion block.

The barrel shifter **816** normalizes any changes to the packet header that occurred during internal routing through the chassis. One function of the internal hardware header of a packet is to permit the CPU to add an encapsulation to a packet. Encapsulation is used by the CPU to complete operations more efficiently by avoiding having to copy the entire packet into CPU memory and then writing the packet back to the buffer pool. Instead, the CPU performs a small modification to the packet header. For example, this might occur when the CPU determines that a packet must be forwarded, but that the CPU must first add data to the header before forwarding can take place. Alternatively, the CPU might also remove data from the header temporarily to assist with forwarding.

During this process, the CPU might move data within the packet header into a non-standard format. For example, the destination address might appear at the wrong location within the packet for transmission over the network. The barrel shifter **816** analyzes the composition of the packet header and shifts the data within the header to normalize it and correct for any CPU modifications that might have occurred. When the barrel shifter **816** completes operations on the packet header, the packet header data is then in a standard format and is passed to the SA substitution and L3 assist block **818** for further processing.

The SA substitution and L3 assist block **818** performs further modifications on the packet header to prepare the packet for transmission over the network. The SA substitution and L3 assist block **818** replaces the MAC address that is required for routing packets. In an Ethernet environment, each packet header contains a destination address and a source address. The source address must be changed on transmit to reflect which port the packet is being broadcast from. The SA substitution and L3 assist block **818** also modifies other Layer 3 header fields as required, such as changing the IPv4/IPX time to live value or the checksum.

The packet is passed to the VLAN insertion block **820** for further processing. VLAN tags that were removed on receipt anywhere in the chassis are stored in the internal hardware header for future use on transmission. The VLAN insertion block **820** takes the internal hardware header information that is passed from the context track and internal header removal block **814** and reintroduces this information into the outgoing packet header as appropriate. This information includes the packet VLAN ID and the Tag Status.



When the outgoing header packet is reassembled for transmission over the network, the packet is stored in the TX FIFO **822** prior to being passed to the XGMAC interface **824**. The TX FIFO **822** enables the VLAN insertion block **820** to begin processing the next packet without having to wait for the XGMAC interface to become available and enables faster operation by the VLAN insertion block **820**.

Additionally, the TX FIFO **822** permits data flow through asynchronous boundaries. In some embodiments of the invention, the TXPP **120** operates at a different speed than the MAC **102**. Data flow must be synchronized between asynchronous components so the TX FIFO **822** acts as a bridge between these components. For example, in the Foundry BigIron switch, the MAC **102** operates at a 156.25 MHz clock and the TXPP operates at only a 66 MHz clock.

FIG. 9 illustrates a block diagram presenting a high-level schematic of the components of an alternative embodiment of the invention that allows for data transfer over two ports at speeds at or in excess of 10 gigabits per second. As shown, the invention comprises a printed circuit board (“PCB”) **902** used to house and provide interconnections for two media access controllers (“MAC”) **904a** and **904b**, a packet processor (“PP”) **906a** and **906b** coupled to each MAC, one or more content addressable memory (“CAM”) controllers **908a** and **908b**, one or more controllers for random access memories containing parameter information (“PRAM”) **910a** and **910b**, transmit and receive quad data rate (“QDR”) random access memory buffers **912a** and **912b**, a transmission manager **914**, and a backplane interface **916**. According to this embodiment, each transmit and receive QDR RAM **912a** and **912b** is depicted as a single module. Those skilled in the art, however, will recognize that multiple QDR RAM or other high speed memory structures may be used to implement distinct transmit and receive memory buffers.

The PCB **902** provides a surface on which to place other components of the invention. The PCB **902**, also known as a “blade” or “module”, can be inserted into a slot on the chassis of a network traffic management device such as a switch or a router. This modular design allows for flexible configurations with different combinations of blades in the various slots of the chassis according to differing network topologies and switching requirements. Furthermore, additional ports for increased network connectivity may be easily added by plugging additional blades into free slots located in the chassis. The aforementioned components that are mounted on the PCB **902** or blade act in concert to allow data packets to be routed between two ports associated with each MAC **904a** and **904b**, as well as other blades connected to the chassis.

Each MAC **904a** and **904b** provides a discrete interface by which data is received and transmitted to and from the network. In one embodiment, such network data comprises Ethernet packets. Each MAC **904a** and **904b** forwards packets that they receive to their respective PP **906a** and **906b** for further processing, and also receive packets for transmission to the network from its associated PP **906a** and **906b**. The MACs **904a** and **904b** perform data conversion required for the packet processors **906a** and **906b** to process packets for routing within the device chassis and for data that the packet processors process to be transmitted to the network. For example, in one embodiment of the invention, each MAC **904a** and **904b** performs data conversions because network data that they receive from their respective frontend ports comprises 32 bit double data rate (“DDR”) data, whereas the associated packet processors **906a** and **906b** each process only 64 bit single data rate (“SDR”) data. Typically, the MACs are responsible for data validity checking as well as data gathering.

Each packet processor **906a** and **906b** is a processor structure, such as an ASIC or FPGA, responsible for receiving packets from an associated MAC **904a** and **904b**, processing the packets for transmission through the device chassis, and processing packets that they receive from the device chassis via the transmission manager **914** intended for transmission over the network. These two functions, while performed by each packet processor **906a** and **906b** for its respective port, are preferably performed simultaneously and in parallel. There are thus two processing pipelines implemented in each of the packet processors **906a** and **906b**: a receive pipeline for processing network packets intended for transmission within the chassis and a transmit pipeline for processing internally routed packets intended for transmission over the network via one of the two ports.

The receive pipeline of each PP **906a** and **906b** is responsible for packet classification, performing CAM and PRAM lookups, generating packet headers for forwarding packets through a chassis, and preparing packet modifications. Each PP **906a** and **906b** receives network packets from its respective MAC **904a** and **904b** in multi-byte bursts based on scheduling priorities determined at a respective MAC **904a** and **904b**. Each PP **906a** and **906b** examines packets and extracts packet forwarding information from the packets such as the destination address (“DA”) of the packet and the source address (“SA”) of the packet. Each PP **906a** and **906b** extracts the type of service (“TOS”), whether the packet has a virtual local area network (“VLAN”) tag, session related data such as in the case of IPv4 or IPX data, and other additional Layer 3 and Layer 4 information useful in routing the packet through the chassis. Each PP **906a** and **906b** also passes the forwarding information that it extracts from the packet header to a CAM processor, **908a** and **908b** respectively, for further processing.

Each of the CAM controllers or processors **908a** and **908b** takes information forwarded by its respective PP **906a** and **906b** and performs a lookup comparing this information to data stored in a local memory of the CAM processor **908a** and **908b**. If the information matches information stored in the local memory of the CAM processor **908a** and **908b**, additional forwarding information regarding disposition of the packet is available in the local memory of the given packet processor’s PRAM processor **910a** and **910b**, which it may retrieve for future incorporation into the packet header. When such successful CAM matches occur, each of the corresponding PRAM processors **910a** and **910b** retrieves additional forwarding information from its local memory for incorporation into the header of the packet. The packet is reformatted with a new internal hardware header for routing the packet within the chassis, which the PP **906a** and **906b** stores in a corresponding transmit and receive QDR RAM **912a** and **912b** for processing by the transmission manager **914**. This internal hardware header is also sometimes referred to as a chassis header.

The transmit and receive QDR RAM **912a** and **912b** acts as a pipeline buffer in the embodiment of the invention depicted in FIG. 9. The transmit and receive QDR RAM **912a** and **912b** enables each PP **906a** and **906b** to store packets that they process and continue processing the next packet without having to wait for the transmission manager **914** to become available, thereby expediting operations of both the packet processors **906a** and **906b** and the transmission manager **914**. Other buffers are used throughout the system and in its various components, as further described herein, to achieve pipelining and faster packet processing in an analogous manner.

While the packet processors **906a** and **906b** handle traffic to and from their respective MACs **904a** and **904b**, as well as



conversions of packet headers between network packet headers and internal chassis packet headers, the transmission manager **914** handles traffic flow to and from the backplane interface **916**. Similar to the packet processors **906a** and **906b**, the transmission manager **914** is a processor chip that implements a quad pipeline architecture: a receive pipeline for network data to be internally routed within the device chassis and a transmit pipeline for internally routed data intended for network transmission. To support two ports on each blade, the transmission manager **914** implements a separate transmit and receive pipeline for each port in each of its transmit cores **914a** and **914b**. These functions, while performed on the same chip, are preferably performed in parallel according to one embodiment of the invention. In one embodiment of the invention, the transmission manager **914** is an FPGA, although use of other processor types such as an ASIC is within the scope of the invention.

Advantageously and as is explained in greater detail herein, the transmission manager **914** implements two transmission cores **914a** and **914b** to handle the transmission of packets to and from both packet processors **906a** and **906b** in a parallel fashion, thereby implementing the receive and transmit pipelines for each port on a single chip. The transmission manager **914** fetches network data intended for routing through the device chassis from each transmit and receive QDR RAM **912a** and **912b** and stores internally routed data intended for network transmission back into the appropriate transmit and receive QDR RAM **912a** and **912b**. The receive pipeline of each transmission core **914a** and **914b** in the transmission manager **914** retrieves data from the transmit and receive QDR RAM **912a** and **912b** according to instructions issued to the transmission manager **914** by each PP **906a** and **906b** associated with a given transmit and receive QDR RAM **912a** and **912b**. The transmission manager **914** determines data transmission priority for the data retrieved and schedules transmissions to the backplane **916** according to this priority scheme through implementation of a plurality of transmission queues or FIFOs, which are used to alleviate head of line blocking issues between packets destined for other blades in the chassis.

The transmit pipeline of each transmission core **914a** and **914b** in the transmission manager **914** handles internally routed packets received from the backplane interface **916** and intended for transmission over the network. The transmission manager **914** collects packets from the backplane interface **916** and organizes them into per-source, per-priority transmit queues stored in the transmit and receive QDR RAM **912a** and **912b**. Each transmit core **914a** and **914b** in the transmission manager **914** notifies its respective PP **906a** and **906b** when a packet is stored in the transmit and receive QDR RAM **912a** and **912b** is available for processing.

One embodiment of the data flow between the packet processors, their associated QDR RAM and each transmit core of the transmission manager is presented in FIG. **10**. Each transmission port (not pictured) on the blade **1000** comprises a connection to one of a plurality of given packet processors **1002a** and **1002b**. As described above and represented in the figures, the packet processors **1002a** and **1002b** provide functionality that provides for initial processing of packets, including CAM and PRAM lookups that are used to modify header information that each packet comprises. By modifying packet header information, the packet processors **1002a** and **1002b** provide an indication as to how the transmission manager **1008** should route the packets to their destination.

Each of the packet processors **1002a** and **1002b** implement multiple data pipelines to speed up data processing. Accordingly, each packet processor **1002a** and **1002b** implements a

receive pipeline **1012** and a transmit pipeline **1014**. When packet processing that the packet processors perform is complete, the packet is put into a receive QDR RAM **1004a** and **1004b**. The receive QDR RAM **1004a** and **1004b** serves as a buffer in the receive pipeline **1012**, allowing the packet processors **1002a** and **1002b** to continue processing incoming packets that the MAC interfaces (not pictured) are forwarding.

The transmission manager **1008** implements a dual transmit core architecture whereby each of the transmit cores **1010a** and **1010b** provides transmission and reception functionality to a given one of the packet processors **1002a** and **1002b**. When the given transmit core **1010a** and **1010b** associated with a given receive QDR RAM **1004a** and **1004b** is available, the transmit core **1010a** and **1010b** gets the next data packet from its receive QDR RAM **1004a** or **1004b** for processing and dispatch, e.g., over the backplane or to an alternate transmit core. Likewise, the transmit core **1010a** and **1010b** processes packets that it receives, e.g., from the backplane or alternate transmit core, and dispatches the packets to the transmit QDR RAM **1006a** and **1006b** over its transmit pipeline **1014**. As can be seen, each packet processor **1002a** and **1002b** implements two data pipelines, one transmit and one receive, whereas the transmission manager **1008** implements four data pipelines, one transmit and one receive pipeline for each transmit core **1010a** and **1010b**.

Receive QDR RAM **1004a** and **1004b**, as well as transmit QDR RAM **1006a** and **1006b**, is particularly suitable for pipelined packet processing and storage in high-speed networking applications where symmetrical high write and read bandwidth is needed. Unlike dualport memory illustrated in conjunction with other embodiments of the invention, QDR RAM has only one address port for both write and read addressing, e.g., write addresses and read addresses must be multiplexed.

While previous embodiments demonstrate the external placement of the multiplexing structure, the packet processor **1002a** and **1002b** and transmission manager **1008** may implement the multiplexer to achieve the goals of higher integration and reducing the number of components. In the receive direction, the packet processor **1002a** and **1002b** writes to the receive QDR RAM **1004a** and **1004b** and the transmission manager **1008** reads from the receive QDR RAM **1004a** and **1004b**. The transmission manager **1008** may implement a receive QDR address multiplexer (not pictured) whereby the packet processor **1002a** and **1002b** sends the receive QDR write address to the transmission manager **1008**. In the transmit direction, the transmission manager **1008** writes to the transmit QDR RAM **1006a** and **1006b** and the packet processor **1002a** and **1002b** reads from the transmit QDR RAM **1004a** and **1004b**. The packet processor **1002a** and **1002b** may implement a transmit QDR address multiplexer (not pictured) whereby the transmission manager sends the transmit QDR write address to the packet processor **1002a** and **1002b**.

FIG. **11** illustrates a more detailed block diagram of the internal architecture of the transmission manager according to one embodiment of the present invention. Each of the ports on a given blade, e.g., port **0** and port **1**, is associated with a given receive QDR RAM **1104a** and **1104b**, respectively. These QDR RAM **1104a** and **1104b**, which act as high speed ram buffers, queue packets that have been processed and are ready for transmission to their destination, e.g., a port on the blade or another blade connected via the backplane to the chassis. According to some embodiments, a priority scheme is implemented in the receive QDR RAM **1104a** and **1104b** to prioritize the sequence in which the transmission manager



**1102** retrieves packets for processing. Within the transmission manager **1102**, a transmission core **1108a** and **1108b** associated with each port retrieves queued packets from its receive QDR RAM **1104a** and **1104b**, respectively, for processing. It should also be noted that, in addition to the components depicted in FIG. **11**, the transmission manager may implement some or all of the structures and functionality associated with other embodiments of the invention described heretofore. Furthermore, it should be apparent to one of skill in the art from the subsequent discussion that the number of transmit cores **1108a** and **1108b** and, therefore, the number of ports that a given blade can support, is limited only by the physical space available on the chip on which the transmission manager **1102** is implemented.

The transmit cores **1108a** and **1108b** fetch data from the receive QDR RAM **1104a** and **1104b** based on instructions received by other components on the blade. At the start-of-packet boundary, each of the transmit cores **1108a** and **1108b** extracts a forwarding identifier (“FID”) from each packet and sends the FID to a FID lookup arbiter **1112**. The FID is an abstract chassis/system wide number used to forward packets. Each packet type has a FID to instruct the blade how to handle a given type of packet. This allows each blade in the chassis to look at the FID separately to decide how to individually forward the packet. For example, the FID can instruct the transmit core **1108a** and **1108b** to transmit a given packet to one or more destinations on the backplane, a port on the blade receiving the packet, a port on the blade other than the one receiving the packet or a combination of these destinations. As will be explained herein, each transmit core **1108a** and **1108b** implements a local switching FIFO **1110a** and **1110b**, respectively, to pass packets between transmit cores without requiring the packet to be transmitted over the backplane, e.g., “one arm routing” between transmission cores.

Because each transmission core is implementing a receive pipeline, FID values for two sets of data must be processed to determine routing information. The FID lookup arbiter **1112** implements an algorithm to equitably select packets for processing. According to one embodiment, the FID lookup arbiter **1112** uses the priority scheme implemented in the receive QDR RAM **1104a** and **1104b** to select packets for FID processing. The FID lookup arbiter **1112** translates the FID of the packet that it selects into a port mask by performing a lookup against addresses stored in the receive FID RAM **1114**. The port mask is a multi-bit field representing the two ports on the blade and other possible backplane slot destinations in the device chassis. FIG. **12** presents one embodiment of a FID port mask. According to this embodiment, the port mask **1202** is a 9-bit field representing the two 10 Gigabit Ethernet ports on the blade and seven other possible backplane slot destinations. By setting one or more bits in the port mask **1202**, each transmission core **1108a** and **1108b** is able to determine the proper destination or destinations for a given packet.

The transmit cores **1108a** and **1108b** utilize the FID port mask to determine how to properly route a packet. To allow packets to be routed between ports on a given blade, each transmit core **1108a** and **1108b** implements a local switching FIFO **1110a** and **1110b**. Where a given transmit core determines that a packet that it is currently processing is destined for its counterpart port, e.g., port **0** receives a packet which transmit core **0** determines as a result of FID processing that the blade must transmit the packet over port **1**, the transmit core places the packet on the local switching FIFO **1110a** and **1110b** of the destination transmit core. Alternatively, each transmit core **1108a** and **1108b** may place packets on its own local switching FIFO **1110a** and **1110b**, respectively, which is polled in accordance with a polling schedule to retrieve pack-

ets for transmission in its associated port. Essentially, the local switching FIFOs **1110a** and **1110b** provide a data pathway for switching packets between ports for transmission without having to route the packets over the backplane, unless a copy of the packet is also destined for a location on the backplane. As is explained herein, packets that the transmit cores **1108a** and **1108b** pull from the local switching FIFO **1110a** and **1110b** for transmission over its associated port are passed to the transmission quality of service module **1128a** and **1128b** for transmission scheduling.

Where the transmit cores **1108a** and **1108b** determine that they must route a packet to the backplane, a receive FIFO arbiter **1118** manages data traffic out of the receive pipeline of each transmit core **1108a** and **1108b**. The receive FIFO arbiter **1118** acts as a buffer in the receive pipeline of the transmission manager **1102** by queuing packets from the receive pipeline of the transmission cores **1108a** and **1108b**. By providing the FIFO buffer, the receive FIFO arbiter **1118** frees the transmit cores **1108a** and **1108b** to continue pulling packets from their respective receive QDR RAM **1104a** and **1104b** for processing. Packets are queued until subsequent components in the receive pipeline **1120**, **1122** are free. The receive FIFO arbiter **1118** analyzes traffic output from the transmit cores **1108a** and **1108b** to calculate from which transmit core to accept packets for placement on the FIFO during a given clock cycle.

The backplane transmit grouper **1120** funnels 32 byte wide data that each transmit core **1108a** and **1108b** fetches from the receive QDR RAM **1104a** and **1104b** into multiple 4 byte wide data that matches the width of the backplane, in addition to generating sideband information for the backplane **1130**. The backplane transmit grouper **1120** is also aware of which backplane destination or destinations to which a given QDR RAM data line wishes to transfer data. In the case where multiple backplane destinations are indicated the backplane transmit grouper **1120** performs data line replication to multiple backplane destinations. The backplane transmit grouper **1120** may also comprise multiple asynchronous FIFOs that couple data from the QDR RAM at a specific clock rate into the backplane clock rate. The backplane transmit grouper **1120** may generate four byte wide data plus sideband information for processing by the backplane transmit sorter **1122**.

When components comprising the receive pipeline of each transmit core **1108a** and **1108b** process packets from the receive QDR RAM **1104a** and **1104b**, the components modify the header information associated with each packet to instruct the transmission manager **1102** how it should route the packet through the pipeline. The backplane transmit sorter **1122** resides at the boundary between the transmission manager **1102** and the backplane **1130**, acting as a regulator to control the flow and destination of data that the transmission manager **1102** processes. The backplane transmit sorter **1122** implements a number of dispatch FIFOs **1132** onto which packets of data are put from the backplane transmit grouper **1120** for transmission over the backplane **1130**. According to one embodiment, the backplane transmit sorter **1122** implements a number of dispatch FIFOs **1132** equal to the number of possible destination slots comprising the chassis that are connected to the backplane **1130**.

Where the backplane transmit sorter implements a lesser number of FIFOs, e.g., one FIFO for all data packets that the transmission manager **1102** is dispatching to the back plane **1130**, “head of line blocking” typically occurs with high frequency. Head of line blocking may be explained as follows: suppose an input link in a switch has an associated FIFO which contains a packet destined for a given output link. Where the given output link is occupied or otherwise busy, the



input link is forced to idle until the output link becomes available, thereby causing other packets in the FIFO destined for other output links to potentially be held up unnecessarily. By applying a sorting algorithm and a number of dispatch FIFOs **1132** approaching the number of destination slots on the backplane **1130**, the blocking condition is significantly reduced or possibly eliminated. According to embodiments of the invention, the backplane transmit sorter **1122** comprises logic that is operative to monitor the availability status of destinations on the backplane **1130** and reorder the packets in each of the dispatch FIFOs **1132** to optimize transmission of the packets over the backplane **1130**. In this regard, a round-robin arbiter may examine backplane destination congestion status as well as FIFO **1132** availability to decide which FIFOs **1132** to read from and in turn transmit data to the backplane **1130**.

As packets arrive at the transmission manager **1102** from the backplane **1130** for transmission over one of the blade's ports, a backplane transmit FID lookup module **1124** processes each data packet. The backplane transmit FID lookup module performs a lookup against destination data in a transmit FID RAM **1126**. The transmit FID RAM **1126** returns a two bit value that instructs the backplane transmit FID lookup module **1124** the appropriate transmit core **1108a** and **1108b** to which the packet is to be sent. Because the transmit FID RAM **1126** returns a two bit value, the backplane transmit FID lookup **1124** may send the packet to transmit core **0** (**1108a**), transmit core **1** (**1108a**), both transmit cores, or neither transmit core, in which case the backplane transmit FID lookup module **1124** drops the packet. Packets can additionally or alternatively be sent to other destinations, such as a CPU for processing or a monitor for debugging purposes. When the backplane transmit FID lookup module **1124** completes processing the data packet, which may comprise modifying the information in the packet's header, it forwards the packet to the proper destination, e.g., transmit core **1108a** and **1108b**.

Each transmit core **1108a** and **1108b** receives data packets from the backplane transmit FID lookup module **1124** and places the packets in a local transmit FIFO **1134a** and **1134b** for storage prior to transmission scheduling by the transmission quality of service (QOS) component **1128a** and **1128b**, respectively. According to some situations, data packets in the local transmit FIFO **1134a** and **1134b** and data packets in the local switching FIFO **1110a** and **1110b** in a given transmit core are ready for transmission at the same time. The transmission QOS module **1128a** and **1128b** may implement one or more QOS techniques that are well known to those of skill in the art, which provide functionality to measure, improve, and possibly guarantee transmission rates, error rates and other characteristics in advance of transmission.

According to one embodiment, the header of each data packet that the transmission manager **1102** receives from the backplane **1130** comprises QOS data. The transmit QOS module examines the QOS header information in each packet at the head of both the local transmit FIFO **1134a** and **1134b** and the local switching FIFO **1110a** and **1110b**. The transmit QOS module **1128a** and **1128b** selects a data packet from one of the FIFOs based on the QOS header information in the data packet and according to a given QOS algorithm that the transmit QOS module **1128a** and **1128b** implements. The data packet that the transmit QOS module selects is placed in an associated transmit QDR RAM **1106a** and **1106b** associated with the transmit pipeline for the given transmit core of which the transmit QOS module is part. The transmit QDR RAM **1106a** and **1106b** buffers packets that a given transmit core **1108a** and **1108b** in the transmission manager **1102**

processes until the packet processor that is associated with the destination port is available to get the next packet in a given transmit QDR RAM **1106a** and **1106b**.

While the invention has been described and illustrated in connection with preferred embodiments, many variations and modifications as will be evident to those skilled in this art may be made without departing from the spirit and scope of the invention, and the invention is thus not to be limited to the precise details of methodology or construction set forth above as such variations and modification are intended to be included within the scope of the invention.

What is claimed is:

1. A system for data transmission comprising:
  - a first integrated circuit (IC) with a plurality of cores, wherein each said core is operative to provide a bi-directional pipeline for packets corresponding to a given one of a plurality of media access control (MAC) interfaces, each said core is further operative to receive packets from and dispatch packets to at least one memory structure, each said core comprises a local switching circuit to transfer packets between the core and at least one other of the cores, and the first IC comprises a backplane transmit sorter circuit that controls dispatch of packets from the plurality of cores to their intended destinations.
  2. The system of claim 1 further comprising:
    - the plurality of MAC interfaces for facilitating receipt and transmission of packets over an associated set of physical interfaces; and
    - second and third ICs coupled to the plurality of MAC interfaces and the at least one memory structure, each of the second and third ICs configured to:
      - perform initial processing of packets received from a respective associated MAC interface of the plurality of MAC interfaces,
      - schedule the transmission of packets to the respective associated MAC interface, and
      - dispatch and retrieve packets to and from the at least one memory structure.
    3. The system of claim 1 wherein the at least one memory structure is a quad data rate (QDR) memory module.
    4. The system of claim 1 wherein the first IC is a field programmable gate array (FPGA).
    5. The system of claim 1 wherein the local switching circuit of each said core of the plurality of cores of the first IC comprises a FIFO.
    6. The system of claim 1 wherein each said core of the plurality of cores of the first IC comprises a transmit quality of service module circuit operative to determine an order in which packets received at the said core from a backplane and the local switching circuit are transmitted to the at least one memory structure.
    7. The system of claim 1 wherein the backplane transmit sorter circuit implements a number of FIFOs to organize packets prior to transmission to one or more slots on a backplane.
    8. The system of claim 7 wherein the backplane transmit sorter circuit implements a number of FIFOs equal to a number of slots on the backplane.
    9. The system of claim 1 wherein the first IC comprises a receive FIFO arbiter circuit operative to arbitrate the receipt of data from the plurality of cores of the first IC for placement in one or more receive FIFOs.



10. The system of claim 1 wherein the first IC comprises a backplane transmit grouper circuit operative to modify a size of a data packet for transmission over a backplane.

\* \* \* \* \*



UNITED STATES PATENT AND TRADEMARK OFFICE  
**CERTIFICATE OF CORRECTION**

PATENT NO. : 8,811,390 B2  
APPLICATION NO. : 12/608972  
DATED : August 19, 2014  
INVENTOR(S) : Yuen Fai Wong

Page 1 of 1

It is certified that error appears in the above-identified patent and that said Letters Patent is hereby corrected as shown below:

On the Title page, in the left-hand column, in item (73) titled ASSIGNEE, delete “Foundry Networks, LLC, Santa Clara, CA (US)” and insert

-- “Brocade Communications Systems, Inc., San Jose, CA (US)” --

Signed and Sealed this  
Twenty-fifth Day of November, 2014



Michelle K. Lee  
*Deputy Director of the United States Patent and Trademark Office*